

Data Dissemination in Unified Dynamic Wireless Networks

Magnús M. Halldórsson* Tigran Tonoyan* Yuexuan Wang[†] Dongxiao Yu[‡]

May 10, 2016

Abstract

We give efficient algorithms for the fundamental problems of Broadcast and Local Broadcast in dynamic wireless networks. We propose a general model of communication which captures and includes both fading models (like SINR) and graph-based models (such as quasi unit disc graphs, bounded-independence graphs, and protocol model). The only requirement is that the nodes can be embedded in a bounded growth quasi-metric, which is the weakest condition known to ensure distributed operability. Both the nodes and the links of the network are dynamic: nodes can come and go, while the signal strength on links can go up or down.

The results improve some of the known bounds even in the static setting, including an optimal algorithm for local broadcasting in the SINR model, which is additionally uniform (independent of network size). An essential component is a procedure for balancing contention, which has potentially wide applicability. The results illustrate the importance of carrier sensing, a stock feature of wireless nodes today, which we encapsulate in primitives to better explore its uses and usefulness.

*ICE-TCS, School of Computer Science, Reykjavik University, Iceland; mmh@ru.is, ttonoyan@gmail.com.

[†]College of Comp. Sci. & Tech., Zhejiang University, Hangzhou, P.R. China; Department of Computer Science, The University of Hong Kong, Hong Kong, P.R. China; amywang@hku.hk.

[‡]Services Computing Technology and System Lab, Cluster and Grid Computing Lab in the School of Computer Science and Technology, Huazhong University of Science and Technology, 1037 Luoyu Road, Wuhan 430074, P.R. China; dxyu@hust.edu.cn.

1 Introduction

Wireless networks are ubiquitous and are on their way to become even more prevalent, e.g., with the advent of Internet-of-Things. Wireless communication is, however, particularly challenging to model algorithmically. In two crucial interrelated aspects, wireless networks on the ground differ from models typically assumed in algorithmic studies. One is the *communication modeling*: when is a transmission successfully decoded, as a function of the environment and the *interference* from other transmissions. The other is *variability with time*: wireless networks are particularly susceptible to changes. Both of these are hard to capture accurately with well-defined, clear-cut rules.

We aim in this paper to address core information dissemination problems – *local broadcast* and (global) *broadcast* – in dynamic distributed networks, under very weak assumptions on the communication. To this end, we propose a communication model with significant flexibility that allows for adversarial control, generalizing essentially all known analytic wireless models. It allows for more general interference relationships than treated before. The network can experience adversarial dynamic behavior, both edge changes (change of signal strengths) and node insertions/deletions.

Wireless communication has traditionally been modeled theoretically by graphs, either geometric or general. Interference is then also transmitted on graph edges (precluding a node from receiving a message from a neighbor if another neighbor is transmitting), but is sometimes represented by a supergraph. *Fading channel* or *physical* models common in communication engineering, on the other hand, consider interference as cumulative, decreasing with distance but adding up. They have been popular in recent algorithmic studies, adding more realism to the formulation. The standard assumption of *geometric signal decay* (that signal decreases inverse polynomially with distance) in the SINR model, the prototypical fading channel, is though equally at odds with experimental evidence. Ultimately, it may prove futile to hope for a clean deterministic model, or even a purely stochastic one, without a significant dose of unpredictability and non-determinism.

Wireless communication is commonly closely linked to mobility, as the transceivers are more often than not on the move. Dynamic changes to reception conditions have many causes other than node mobility, since almost any changes in the environment affect transmissions due to reflections of signals over multiple paths, antenna characteristics, scattering, and diffraction. These changes are by nature hard to predict, even when assuming a “mobility model”. The most robust approach would then be to assume a non-trivial adversarial component. While the study of algorithms in dynamic networks has a long history, little has been done in cases where interference plays a role.

Setting and Model. Nodes are distributed and autonomous. There is no built-in structure and the nodes have no information besides bounds on model parameters and an upper bound on the number of nodes. Communication is locally synchronous, but there is no global clock.

We will assume a very general model for when nodes successfully communicate. Nodes are located in space, with separation between points given by the relative decrease in strength of signal (or interference) sent between the points. This induces a metric space, when these decays are raised to the appropriate power – actually, it is a *quasi-metric*, since symmetry need not hold. For distributed computation to be possible, the quasi-metric must have *bounded independence* (to be defined precisely). *Edge changes* can occur, with some restrictions, which are changes in the signal strength between the pair of points.

The rule for when communication is successful is only partially pre-specified: transmission succeeds on a “clear channel”. That is, if a sender is within a *communication radius* from the receiver, if no other node transmits within a (larger) radius, and if the combined interference from all other transmitting nodes is (quite) small, then the transmission succeeds. Otherwise, success is up to the adversary, or it can be further specified by the particular model assumptions desired.

This captures essentially all known algorithmic wireless models (including quasi-unit disc graphs, unit-ball graphs, bounded-independence graphs, k -hop extensions, and SINR). The only exception is the radio network model with general graphs, which cannot be extended to involve comprehensive interference without a major hit in time complexity.

The generality of our model is a key feature. Given the vagaries of actual wireless environments, it is preferable for robustness reasons to make minimal assumptions about the communication model. A conservative approach is then to seek algorithms that work in most established models rather than depending on model-specific factors that simplify the life of the algorithm designer.

Our Approach and Results. The key algorithmic technique is a natural randomized *contention balancing* procedure, where a node continuously adjusts its transmission probability based on the interference that it senses. It allows nodes to stabilize quickly from any initial conditions, or after waking up. This routine is a variation on an old story, a simple backoff procedure to manage local contention:

if a node v senses contention in a given round beyond a fixed threshold, then v halves its transmission probability in the next round and otherwise doubles it.

Our main technical contribution is to show logarithmic-round convergence of this method to a steady state of nearly balanced contention, from an arbitrary starting configuration and in the presence of network changes. It proves also to be surprisingly tolerant of different communication models. The higher level algorithms are then built on top of this primitive.

A crucial component is the use of *carrier sense* to detect the cumulative amount of signals in the air. Since it is supplied by the cheapest available hardware today as RSS (received signal strength) readings, we posit that carrier-sense capability should be the default assumption in wireless algorithmics (while exploring the necessity of different assumptions is interesting theoretically). As carrier-sense indicators can provide fine-grained information, we are interested in restraining its use and identifying which aspects are necessary to achieve the results obtained. To this end, we identify several *primitives* that carrier-sense can supply, restrict the algorithm to use only a subset of the primitives, and examine which of these are truly necessary.

The local broadcast algorithm simply runs the contention balancing procedure, with nodes bowing out when they are sure to have completed their transmission. The broadcast algorithms are based on sparsifying the instance, so that only nodes of constant density actually participate in the global broadcast action. The former is achieved in $O(\Delta + \log n)$ time, where Δ is the maximum number of neighbors that a node can have, while the latter takes $O(\hat{D})$ rounds, where \hat{D} is a dynamic diameter.

These dissemination algorithms are efficient enough to improve on some of the results known for static versions of the problems. The local broadcast algorithm is strongly optimal, or within constant factors on every instance. In the standard setting (static, spontaneous case), the algorithm is *uniform*, in that it need not know the network size. The broadcast algorithm is also optimal and uniform in the same setting, while in the non-spontaneous setting it is faster by a logarithmic factor than the previous algorithm of [33] that however does not require carrier sense.

Closely Related Work. There are two largely disjoint bodies of work of wireless algorithmic results, with work on fading models like SINR slowly catching up with the better studied graph-based models. One approach for capturing more realism in SINR model is to move beyond Euclidean metrics [15], even to general ones [36]. One can view relative signal decrease as implicitly defining a quasi-distance metric [5]. Link scheduling problems can be formulated on edge-weighted interference graphs [30] that properly generalize both graph-based and SINR models, linked by a graph-theoretic parameter. Distributed dissemination problems, however, necessarily require metric restrictions,

such as doubling or *fading* metrics [23], and limits on communication abilities in order to capture both types of models.

The local broadcast and global broadcast problems have been extensively studied in both graph-based radio network models [1, 2, 11, 12, 13, 17, 18, 37, 43, 46] and the SINR model [4, 12, 21, 28, 33, 47, 51, 52, 54]. For local broadcast, the best results known in the radio network model are both $O(\Delta \log n)$ with and without knowing (an upper bound on) Δ [1, 18]. In the SINR model, with knowledge of Δ , the local broadcast can be accomplished in the same time bound as in the radio network model [21]. If Δ is not known, the best result is $O(\Delta \log n + \log^2 n)$ [28, 51], which is improved to $O(\Delta + \log n)$ with free acknowledgments [28]. This can be further improved to $O(\Delta + \log n \cdot \log \log n)$ in the spontaneous case, when Δ is known [4].

The time complexity of non-spontaneous broadcasting in the radio network model is $\Theta(D + \log n) \log(n/D)$ [2, 11, 37, 43, 46] without collision detection. With collision detection, this lower bound was recently broken in [17], where a solution of $O(D + \text{polylog}(n))$ was given. Broadcasting has also been treated in the SINR model under a variety of assumptions. Some are stronger than ours (location information [32, 34], power control [52]), while others relax the assumption about the connectivity property, incurring necessarily much higher complexity [8, 12, 35]. Results in our setting, but without carrier sensing, include time complexity of $O(D \log n \cdot \text{polylog}(R_s))$ [12] (see also [24]), where R_s denotes the maximum ratio between distances of stations connected in the communication graph; and $O(D \log^2 n)$ [33]. In the spontaneous setting, where the nodes can build an overlay structure along which the message is then propagated, Scheideler et al. [47] used carrier sense to give a dominator algorithms, which can be applied to solve broadcast in $O(D + \log n)$ rounds. Yu et al. [52] solved the problem in the same time bound using power control, while the algorithm in [33] that requires neither power control nor carrier sense runs in time $O(D \log n + \log^2 n)$.

These problems have also been treated in dynamic networks. In the *unstructured* model [40], where nodes may wake up asynchronously (modeling the node insertion), the local broadcast problem is well studied, even in the SINR setting [21, 28, 51, 54], but this model does not consider node deletion. In the *dual graph* model [38] (originally due to [10]), both the local broadcast [18] and global broadcast [19, 20, 38, 39] problems are studied. But this model involves only edge behavior and not node changes (churn). Hence, the impact of dynamicity on wireless information dissemination is still largely unexplored.

More detailed related work is introduced in Sec. A.

Our Contributions. We have obtained generalized and improved algorithms for two of the most fundamental dissemination problems, in some cases improving the best results known in static settings. Beyond these specific results, we identify the following technical contributions:

1. *Unified model of wireless networks.* The model proposed appears to be the first that allows for the development of pan-model distributed dissemination algorithms. This hopefully prompts further studies crossing the artificial boundary between graph- and fading-based models.
2. *Dynamic networks under interference.* This appears to be the first work to address dynamic networks in the presence of comprehensive interference.
3. *Uniform algorithms.* Our algorithms in the static spontaneous setting appear to be the first in fading models that work independent of instance parameters (number of nodes, max. degree).
4. *Primitives for carrier-sense.* We introduce several primitives or *capabilities* that can be implemented using environmental sensing, and propose to study the power of such primitives.
5. *Stabilization mechanisms.* We identify contention adaptation as a fundamental ability in wireless networks, that appears to be of crucial value to implement other distributed tasks.

Roadmap. The formal model and basic definitions, including the definitions of communication model and carrier sensing primitives, are given in Sec. 2. Sec. 3 contains the core technical part

of the paper that includes the basic contention balance routine and its analysis. The main results concerning local and global broadcast problems are presented in Sections 4 and 5, respectively. Due to space constraints, most proofs are relegated to appendices.

2 Models and Definitions

We consider a dynamic network of point-size wireless devices (nodes). Nodes can transmit messages in time slots/rounds that are sufficiently long to allow a transmission of a single message. No global clock or synchronization of rounds is required, but the clocks of different nodes run at a similar rate, i.e., the length of a round differs between nodes at most by a factor of 2. Nodes may arrive and leave the network at any time. Unless specified otherwise, the nodes are assumed to work *non-spontaneously*: they can initially be in sleep state and join the execution of an algorithm only after receiving a message. We say a node is *alive* at some point in time if it is present in the network. We assume the total number of nodes in the network is polynomially bounded by a number n in each round. We use V to denote the set of alive nodes at any fixed point in time and also use n to denote the current number of nodes, i.e. $n = |V|$.

We assume all nodes use the same transmission power P for communication in all rounds.

Metrics. The *signal strength* – or *interference*, depending on context – of transmitting node u on a node v is $\mathcal{I}_{uv} = P/f(u, v)$, where $f(u, v) > 0$ is the *path loss* from u to v . The *metricity* of a space (V, f) is the smallest number ζ such that for every triplet $u, v, w \in V$, $f(u, v)^{1/\zeta} \leq f(u, w)^{1/\zeta} + f(w, v)^{1/\zeta}$ [5]. We define $d(u, v) = f(u, v)^{1/\zeta}$ if u and v are different nodes and $d(u, v) = 0$ when $u = v$. Note that (V, d) is a *quasi-metric*, as all metric axioms except symmetry hold. In the rest of the paper, we assume that in each round the metricity of the network is bounded by a fixed constant ζ and will work with values $d(u, v)$ instead of $f(u, v)$. We assume the quasi-metric (V, d) has *bounded independence*, defined below, roughly stating that there cannot be many nodes each causing high interference to a *fixed* node, while having low mutual interferences.

First, some notations. The *ball* with radius r centered at u is defined as $B(u, r) = \{v \in V \mid \max\{d(v, u), d(u, v)\} < r\}$. The *in-ball* with radius r centered at u is defined as $D(u, r) = \{v \in V \mid d(v, u) < r\}$; clearly, $B(u, r) \subseteq D(u, r)$. A set $S \subseteq V$ is a *r -packing* for set S' if balls of radius r centered at nodes in S are contained in S' and are disjoint. S is a *r' -cover* for S' if the union of balls of radius r' centered at nodes in S contains S' . Note that any maximal r -packing is a $2r$ -cover, and thus one can bound sizes of covers by packings.

We say that (V, d) has (r_{\min}, λ) -*bounded independence*, for given $r_{\min} \geq 0$ and $\lambda > 0$, if for every $q \geq 1$ and every in-ball D of radius qr_{\min} , the size of a maximum cardinality r_{\min} -packing of D is at most $C \cdot q^\lambda$, where C is a constant, possibly depending on λ . For instance, the Euclidean plane is $(r, \lambda = 2)$ -bounded independent, for every r .

Neighborhoods, Communication Graph and Dissemination Problems. Let R denote the maximum transmission distance possible when no other node transmits. As the latter event is arguably very rare, we define the *communication radius* $R_B = (1 - \epsilon)R$ as a slightly smaller distance, where ϵ is a *precision* parameter. We will drop the parameter ϵ whenever it is fixed and clear from the context. Fix a round t . The *neighborhood* of a node is $N_t(u, \epsilon) = \{v \in V : d(u, v) \leq (1 - \epsilon)R\}$, describing who u can communicate with directly. The basic operation of interest is when u broadcasts a message to its neighbors $N_t(u)$. The *communication graph* is a directed graph $G_t(V, E)$, where $(u, v) \in E$ if and only if $v \in N_t(u)$. Thus, the sequence G_0, G_1, \dots defines a *dynamic graph*. The *vicinity* of u refers to a larger region, $D_u^\rho = D(u, \rho R)$, for a parameter $\rho > 1$.

The data dissemination problems that we consider are defined below. We say that a node u *mass-delivers* in round t if it transmits and all its neighbors ($N_t(u)$) receive the message.

- In the *Local Broadcast* problem, given a node u , it is required to minimize the time from the beginning of the algorithm until node u mass-delivers at least once, assuming it stays alive during that time.
- In the (*global*) *Broadcast* problem, given a distinguished *source* node that initially holds a message, the goal is to minimize the time needed to deliver the message to every node in the network through multihop transmissions.

One Hop Communication. When is a transmission successfully received? Suppose a node u transmits in round t , and let S be the set of concurrently transmitting nodes. Let $\rho_c = \rho_c(\epsilon) \geq 0$ and $\mathcal{I}_c = \mathcal{I}_c(\epsilon) > 0$ be parameters that depend on the precision ϵ . ρ_c
 \mathcal{I}_c

Definition 1. (*SuccClear, Success on a clear channel*) If no other node in $D_u^{\rho_c}$ transmits and the total interference at node u is at most \mathcal{I}_c (i.e. $S \cap D_u^{\rho_c} = \emptyset$ and $\sum_{v \in S} \mathcal{I}_{vu} \leq \mathcal{I}_c$), then the transmission of node u is successfully received by all its neighbors ($N_t(u)$). Otherwise, the reception is under adversarial control.

Randomized Algorithms. We mainly consider randomized algorithms of the following form: in each round t , node v makes a transmission with probability $p_t(v)$, independent of other nodes' transmissions in that round. An important notion for the analysis of such algorithms is *local contention*, the sum of the transmission probabilities in a close region. The contention in the *close neighborhood* of a node v in round t is $P_t(v) = \sum_{w \in B(v, R/2)} p_t(w)$, where the radius $R/2$ allows all pair of nodes in $B(v, R/2)$ to potentially communicate. Also, let $P_t^\rho(v) = \sum_{u \in D_v^\rho} p_t(u)$ denote the contention in the larger *vicinity* of v in round t (ρ will be fixed later). We will also use the notation $\mathcal{I}_t^\rho(v)$ to denote the interference at v from nodes outside its vicinity (in $\bar{D}_v^\rho = V \setminus D_v^\rho$) in round t . The expected value of $\mathcal{I}_t^\rho(v)$ is $\hat{\mathcal{I}}_t^\rho(v) = \sum_{w \in \bar{D}_v^\rho} p_t(w) \mathcal{I}_{wv}$. $P_t(v)$
 $P_t^\rho(v)$
 $\mathcal{I}_t^\rho(v)$
 $\hat{\mathcal{I}}_t^\rho(v)$

Sensing Primitives. We assume the nodes have abilities to sense activity on the channel. Namely, we assume the nodes are able to detect high and low contention in their vicinity, detect (under some conditions) whether their transmission in a given round succeeded and detect a single very near transmission. In the following, we formalize these notions in three primitives: CD, ACK and NTD. We show in Sec. B how *all* these primitives can be implemented with *basic* physical carrier sensing and possibly also with other means.

Contention Detection (CD). Contention can be probabilistically deduced from measured level of radio activity. We want to relax this ability and will use the following variant, where the outcome of CD is one of the two values: BUSY or IDLE channel. Formally, for each node v and round t :

- if contention among close neighbors is high ($P_t(v) > \phi$) then they *all* detect BUSY channel in round t with probability at least $1 - h_1^{-\phi}$, for given $\phi > 1$, where $h_1 > 1$ is a constant, h_1
- if the contention in the vicinity of v is low ($P_t^\rho(v) \leq \eta$) and the interference on v from outside its vicinity is above a threshold ($\mathcal{I}_t^\rho(v) < \mathcal{I}_{cd}$) then v detects IDLE channel in round t with probability at least $h_2^{-\eta}$, for given $\eta > 0$, where $\mathcal{I}_{cd} > 0$ and $h_2 > 1$ are constants. h_2

Successful Transmission Detection (ACK). If a node u has the ACK=ACK(ϵ) primitive (depending on the precision parameter ϵ) then: if u transmits in round t , the interference at u is bounded by \mathcal{I}_{ack} and the transmission is received by all nodes in $N_t(u, \epsilon)$, then the outcome of ACK is 1, where \mathcal{I}_{ack} is a parameter. If the transmission is not received by a node $v \in N_t(u, \epsilon)$ then the outcome is 0. Otherwise, the outcome is 0 or 1, adversarially. \mathcal{I}_{cd}
 \mathcal{I}_{ack}

Near Transmission Detection (NTD). With NTD=NTD(ϵ) primitive, a node is able to detect if a transmitter is very close, assuming that it receives the transmitted message. The outcome of NTD is 1 for node v in round t if v receives a transmission from a node u , $u \in D_v^{\epsilon/2}$. Otherwise, the outcome of NTD is 0. This can also be made approximate.

Dynamicity. We consider a dynamic network where the topology may change adversarially in each round due to node churn (node arrivals/departures) and edge changes. We assume that arriving nodes start running the algorithms from an initial configuration, so we do not limit the rate of churn. With edge changes, existing nodes that were not neighbors before, may become neighbors (e.g. due to mobility). The new neighbors may cause too much interference in a too short time, so the edge changes should be limited. We assume the amount of edge changes is bounded for each node v , as follows. Consider a time interval T of length $\Omega(\log n)$. We require that the number of new neighbors of v during T (not counting churn) is bounded by $\tau|T|$, where $|T|$ denotes the number of rounds in T and τ is a constant, to be fixed later. We further assume the fraction of rounds in T when there are more than ϕ new neighbors of v is bounded by $O(\phi^{-k})$ for every $\phi \geq 1$, where $k > 2\lambda/(\zeta - \lambda)$. Note that there is no restriction on distance changes inside the neighborhood of v (e.g. it is fine for node v if its neighbors move, as far as they remain neighbors). Note that the edge changes may affect the underlying metric, but we require that the upper bounds on metricity and independence are maintained. τ
 k

Requirements and Assumptions. For the convenience of the reader, we gather together all of our assumptions and requirements in a single place.

Communication is assumed to succeed in a clear channel (**SuccClear**). We assume constant metricity ζ and that (V, d) has (r_{\min}, λ) -bounded independence with $1 \leq \lambda < \zeta$. For the local broadcast problem, we assume that $r_{\min} \leq R/4$, and for the broadcast problem that $r_{\min} \leq \epsilon R/4$. As is standard in fading models, the communication radius R_B is necessarily an $(1 - \epsilon)$ -fraction of the maximum transmission distance in a clear channel.

Besides the knowledge required by the primitives that are needed for a particular algorithm, the nodes are assumed to know the precision parameter ϵ . A polynomial estimate on the number of nodes, n , is needed in dynamic and non-spontaneous algorithms, but not in the static spontaneous problems. Knowledge of approximations of model parameters are needed to implement primitives, including ϵ , ζ , R , ρ_c , and \mathcal{I}_c . Knowledge of the maximum degree Δ is not needed.

Synchronous operation is only assumed in the Broadcast algorithm. Aspects not defined or constrained are assumed to be under (adaptive) adversarial control, including when transmissions that fail **SuccClear** are successful, or when nodes appear or disappear from the network.

The extent of increases in edge strengths over a period is restricted, as detailed above, while decreases are not and neither are node changes. The dissemination problems are only expected to function with the set of nodes that are sufficiently stable, as detailed in the respective section.

3 Controlling Contention

In order to keep the contention in the network balanced, we propose a basic procedure called **TRY&ADJUST**, which will be the main building block in our algorithms. The idea is to let each node adapt its transmission probability to the contention detected using the assumed CD primitive. The parameter $\beta \geq 1$ describes the passiveness of the newly arriving nodes.

- TRY&ADJUST(β):** Each node v maintains transmission probability $p_t(v) \leq 1/2$ in each round t , initialized as $p_t(v) = \frac{1}{2}n^{-\beta}$ when v enters the network. In round $t + 1$, v does:
1. Transmit with probability $p_t(v)$, and
 2. Set $p_{t+1}(v) = \begin{cases} \max\{p_t(v)/2, n^{-\beta}\} & \text{if BUSY channel, and} \\ \min\{2p_t(v), 1/2\}, & \text{otherwise.} \end{cases}$

The aim for controlling contention is, of course, to ensure that transmissions made have a fair chance of being successful, which means they sufficiently overpower the interference experienced at

intended receiver from all other transmissions made in that round. We account for this interference in two ways: the *contention* captures the expected interference from the nodes' neighbors, while the *interference* integrates also the interference from nodes further away.

We will measure the contention in the *vicinity* of each node v , i.e. in D_v^ρ , where ρ is a large enough constant. We specify a threshold $\eta = \log_{h_2}(10/9)$ (recall h_2 from CD definition) for measuring contention: if $P_t^\rho(v) > \eta$ then round t is a *high contention round* for node v and is a *low contention round*, otherwise. We further specify a threshold $\hat{\mathcal{I}}$ for interference: if $\hat{\mathcal{I}}_t^\rho(v) > \hat{\mathcal{I}}$ then round t is a *high interference round* for v and is *low interference round*, otherwise.

These thresholds are chosen so as to ensure that in a low-contention/interference round, node v will be likely to succeed if it transmits. However, requiring all or most rounds to be low contention for all nodes will lead to high delays. Instead, it turns out that most rounds will have *bounded contention* and low interference, which allows for good progress; we say node v experiences *bounded contention* in round t if $P_t^\rho(v) < \hat{\eta}$, where $\hat{\eta} > 0$ is a large enough constant, to be specified later. We say that a round t is *good* for node v if t is both bounded contention and low interference round (in which case, some node in v 's vicinity has a good chance of successfully transmitting).

We analyze the properties of TRY&ADJUST using the notion of a *phase*, the shortest time in which at least $\gamma \log n$ rounds occur for all nodes, where γ is sufficiently large (given in Prop. 3.1). We use H to denote a general phase and also the set of rounds in that phase. $|H|$ denotes the number of rounds in a phase H (i.e. $\gamma \log n$). The fundamental property of TRY&ADJUST (Prop. 3.1) is that, for each node and each phase, most of the rounds in the phase are good for that node. This property is then used to show that:

1. If most of the good rounds in a phase have low contention, then node v detects IDLE channel in most of those rounds.
2. Otherwise, during a constant fraction of the rounds, a node in the vicinity of v mass-delivers.

By choosing the parameters carefully, i.e. requiring *low enough* contention, we can make sure that during a phase with mostly low contention, the node detects IDLE channel in *most of the rounds* (more than half) in a phase, thus leading to an increase of transmission probability by the end of the phase, which, after sufficiently many phases ensures message delivery, w.h.p. On the other hand, during a phase with mostly high contention, there will be many nodes in the vicinity of v that successfully transmit, leading to lower contention. These ideas are applied in Thms. 4.1 and 5.1. The core idea behind this analysis is based on [47].

Proposition 3.1. *Let $\sigma \in (0, 1)$. If constants $\rho = \rho(\sigma, \hat{\mathcal{I}})$, $\hat{\eta} = \hat{\eta}(\rho, \sigma)$ and $\gamma = \gamma(\rho, \hat{\eta}, \sigma, \hat{\mathcal{I}})$ are large enough then for each node v and phase H , with probability $1 - O(n^{-3})$, a $(1 - \sigma)$ -fraction of the rounds in H are good.*

The proof is rather technical and is deferred to Sec. C but the intuition is as follows. The contention in each neighborhood must be bounded most of the time, because when it becomes large, it has a high chance of being decreased due to BUSY channel. Moreover, we show that in expectation, the contention in *all* local neighborhoods is bounded, which is then combined with a geometric argument to show that the expected interference at each node is low most of the time.

We derive from the fundamental property two useful propositions. The first says that if contention is high, then nodes in the vicinity deliver the message.

Proposition 3.2. *Assume that constants $\hat{\eta}, \rho, \gamma$ are large enough. For each node v and phase H , if at least $1/10$ -fraction of the rounds of H are of high contention, then $\Omega(|H|)$ nodes in D_v^ρ mass-deliver, with probability $1 - O(n^{-3})$.*

To this end, we first show that if a round is good for node u and a node in its vicinity transmits, then it mass-delivers with constant probability, utilizing both metric assumptions and the properties

of good rounds. We then argue that since most rounds are good (by Prop. 3.1) and most rounds have by assumption sufficient contention, many rounds will be both good and with sufficient contention, and in each of those, a node in the vicinity of v is likely to transmit and succeed.

When contention is low in a lot of rounds of a phase, the node will detect IDLE channel by the CD primitive in many rounds. This will actually happen during many good rounds, which have the low local contention and low external interference to allow for this detection.

Proposition 3.3. *Assume that $\hat{\eta}, \rho, \gamma$ are large enough. For each node v and phase H , if at least 9/10-fraction of the rounds of H are low contention rounds, then with probability $1 - O(n^{-3})$, in at least 3/5-fraction of the rounds of H , v will detect IDLE channel and have low contention and low interference.*

4 Local Broadcast

We propose an algorithm for asynchronous local broadcast in dynamic networks. The algorithm is an extension of the TRY&ADJUST procedure, where the nodes try to balance the contention in the network and stop transmitting as soon as they deliver their message. We assume the nodes are powered with CD and ACK primitives. Note that the passiveness parameter is set to $\beta = 1$, which means that the transmission probability of nodes does not get below $1/(2n)$.

LOCALBCAST: Each node v executes TRY&ADJUST(1) with the following additional step: if v transmits and detects ACK, it stops (i.e. $p_r(v) = 0$ for $r > t$).

We will estimate the performance of the algorithm using the notion of *dynamic degree*, defined as follows. Given a parameter $\rho > 0$, we denote $\Delta_v^\rho(t, t') = |\cup_{r=t}^{t'} D_v^\rho(r)|$ for node v and rounds t, t' with $t' > t$, where $D_v^\rho(r)$ denotes the in-ball D_v^ρ in round r .

Below we prove that if there are not too many node insertions in the neighborhood of a node v , then v mass-delivers (delivers to all its neighbors) in time comparable to its dynamic degree with ρ a constant. The main tools for proving the bound are Props. 3.2 and 3.3. First we argue that if there is a phase of mostly low contention, then node v will deliver its message, w.h.p. Then we show that if the insertions are not too intensive then the contention around v will decrease and a phase with mostly low contention will happen.

Theorem 4.1. *There is a constant $\rho > 0$, such that a node v performing LOCALBCAST asynchronously in a time interval $T = [t, t']$ with $t' - t = \Omega(\Delta_v^\rho(t, t') + \log n)$ mass-delivers, w.h.p., provided that $t' - t = O(n^2)$.*

Note that the assumption $t' - t = O(n^2)$ is needed only for making the claim w.h.p.: it can be relaxed to higher degree polynomials by only increasing constant factors.

Proof. Let us fix constants $\rho, \hat{\eta}, \gamma$ so that Props. 3.2 and 3.3 with $\hat{\mathcal{I}} = \min\{(1-1/\rho)^\zeta \mathcal{I}_c, \mathcal{I}_{cd}, \mathcal{I}_{ack}\}/10$.

We partition T into phases (for node v) and classify them into two types: (type A) phases H where at least 1/10-fraction of rounds are high contention rounds (i.e. $P_t^\rho(v) \geq \eta$), and (type B) phases H where at least 9/10-fraction of rounds are low contention rounds (i.e. $P_t^\rho(v) < \eta$).

Claim 4.2. *Node v mass-delivers in a type B phase, w.h.p.*

Proof. Let H' be the low contention and interference rounds during phase H where v detects IDLE channel. By Prop. 3.3 $\Pr[|H'| \geq 3|H|/5] > 1 - O(n^{-3})$. Assume for now that the latter happens. For each $t \in H'$ we have $p_{t+1}(v) = \min\{2p_t(v), 1/2\}$. Let us call this operation doubling.

The value of $p_t(v)$ at the beginning of the phase is at least $1/(2n)$, so $\log n$ doubling operations are sufficient to raise it to $1/2$. The probability can be further halved during the phase at most $|H \setminus H'| \leq 2|H|/5$ times. Thus, we may assume we have at most $2|H|/5 + \log n$ halving and at least $3|H|/5$ doubling operations applied to an initial value $1/2$. If $\gamma > 10$, then the total number of halving operations is less than $|H|/2$. It follows that v has $p_t(v) = 1/2$ in at least $(3/5 - 1/2)|H| = |H|/10$ low contention/interference rounds. By Lemma D.1, in each such round, v mass-delivers with probability at least $0.9 \cdot \frac{1}{2} \cdot 4^{-\eta}$; hence, if γ is large enough, v mass-delivers in H , w.h.p. \square

It remains to argue that there will be a type B phase during time interval T . Consider a type A phase H . Prop. 3.2 implies that with probability $1 - O(n^{-3})$, $\Omega(\log n)$ nodes in D_v^ρ deliver their message and stop during phase H . Thus, with probability $1 - O(n^{-1})$, there are at most $C \cdot \frac{\Delta_v^\rho(t, t')}{\log n}$ type A phases with $C > 0$ a constant, as there are at most $\Delta_v^\rho(t, t')$ nodes in D_v^ρ during the time interval T (also recall that $t' - t = O(n^2)$). We conclude that if T consists of at least $C \cdot \frac{\Delta_v^\rho(t, t')}{\log n} + 1$ phases, it will contain a type B phase and v will deliver its message w.h.p. \square

Implications for Static Networks. In static networks, the parameter $\Delta_v^\rho(t, t')$ is at most $|D_v^\rho(0)| = O(\Delta)$ if ρ is constant, where $\Delta = \max_v \{|N(v)|\}$ is the maximum size of a neighborhood in the network. Thus, we obtain the following optimal result (up to constant factors) for static networks, as Δ and $\log n$ are lower bounds even when running in the spontaneous mode [51].

Corollary 4.3. *When running LOCALBCAST in a static asynchronous network, each node v completes local broadcast in $O(|D_v^\rho| + \log n) = O(\Delta + \log n)$ rounds, w.h.p.*

Remark. In the special case when the nodes can start executing the algorithm simultaneously, i.e. in the *spontaneous mode*, the nodes need not know an upper bound on the size of the network. Indeed, each node may start running TRY&ADJUST with initial probability set to an arbitrary value and with no lower limit. The first phase will be spent for stabilization and can be ignored, while the argument for the rest of the phases is nearly identical to the one in Thm. 4.1.

5 Broadcast

For the broadcast problem, we assume nodes communicate in synchronized rounds of equal length. Each round consists of two slots. The idea is to use the first slot of each round for disseminating the message with TRY&ADJUST and the second slot for notifying nodes which have no uninformed neighbors. The latter is accomplished by using higher precision primitives, namely ACK($\epsilon/2$) and SuccClear($\epsilon/2$), when executing TRY&ADJUST. This helps to detect a transmission that is successfully received by all nodes in $N(v, \epsilon/2)$ of a node v . Upon detecting such a transmission, node v resends the message in the second slot, in order to inform nodes u with $v \in D_u^{\epsilon/2}$ that their neighborhood ($N(u, \epsilon)$) has been covered. A node u can detect that $v \in D_u^{\epsilon/2}$ using NTD.

The algorithm is presented below. We assume the passiveness parameter β of TRY&ADJUST is large enough, to be defined later. Note that the algorithm works for the non-spontaneous mode, as nodes act only after receiving the message.

BCAST(β): Initially, only a source node s has the message. A node v , upon receiving a message, starts executing TRY&ADJUST(β) in the first slot of rounds. In addition, in each round t ,

1. if v detects ACK in the first slot, it retransmits in the second slot and restarts TRY&ADJUST(β),
2. if v receives a message in the first slot and detects NTD in the second slot, it restarts TRY&ADJUST(β).

In order to evaluate the progress of the algorithm, we use a notion of a *dynamic distance*, as defined below. Let $c > 0$ be a parameter. A sequence $v_1 = s, v_2, \dots, v_k = v$ is called a *stable s - v path* if there is a sequence I_1, I_2, \dots, I_{k-1} of time intervals with $I_i = [b_i, e_i]$, such that $e_i - b_i \geq c \log n$, $e_i - e_{i-1} \geq c \log n$ and nodes v_{i-1} and v_i are both alive and $v_i \in N(v_{i-1}, \epsilon)$ during I_i . The *time-length* of a stable s - v path is $e_{k-1} - b_1$. The *stable s - v distance* $D_{st}^c(s, v)$ is defined as the minimum time-length of a stable s - v path. Note that a stable path need not be connected at any fixed point in time. Moreover, most of the nodes might be missing at any given point in time. $D_t^c(s, v)$

The core idea behind the analysis of the following theorem is similar to the case of local broadcast: we show that as soon as a neighbor u of a node v has the message and u and v keep being neighbors for $O(\log n)$ rounds, v will receive a transmission of u during those rounds.

Theorem 5.1. *Assume the edge change rate τ is sufficiently small. There are constants $\beta, c > 0$, such that when running $\text{BCAST}(\beta)$ in the synchronous mode, each node v receives the message in $O(D_{st}^c(s, v))$ rounds w.h.p.*

Implications for Static Networks. When the network is static, Theorem 5.1 can be reformulated in terms of hop-distance $\text{dist}_G(s, v)$ in the communication graph, which is defined as the length of the shortest directed s - v path in G : we have that $D_{st}^c(s, v) = O(\log n) \cdot \text{dist}_G(s, v)$ for any node v . Note also that in this setting nodes that succeeded transmitting or detected NTD need not continue the algorithm, so they stop transmitting. In this case, setting the passiveness parameter to $\beta = 1$ suffices. We call this variant of the algorithm BCAST^* .

Corollary 5.2. *When running BCAST^* in synchronous non-spontaneous mode in a static network with source s , each node v receives the message in $O(\log n) \cdot \text{dist}_G(s, v)$ rounds w.h.p. When the communication graph is strongly connected, the broadcast from any source node is completed in $O(\log n) \cdot D_G$ rounds, where D_G is the diameter of the communication graph.*

In the spontaneous mode, the bound above can be further improved to $O(D_G + \log n)$; see Appendix G. This is based on finding a constant-density dominating set in $O(\log n)$ time [47] and simultaneously propagating along the dominators in $O(D_G + \log n)$ time. We can therefore extend the approach based on [47] to uniform algorithms in bounded-independence metrics.

These results are close to best possible. We show below that in order to obtain bounds of that magnitude, it is necessary to have the NTD primitive. To this end, we extend the lower bound construction of [12, Thm. 7] for “compact SINR” to our setting.

This construction leverages the property of our model that there can be arbitrarily many nodes that are mutually close to each other. Namely, the bounded-independence metric is strictly more relaxed than the standard Euclidean metrics. Indeed, there is a $O(D \log^2 n)$ -round broadcast algorithm for the SINR model that does not need NTD or other carrier sensing primitives [33]. What the lower bound then illustrates is that to obtain such results, one must depend on opportune traits of the SINR model that we have tried to avoid and are not necessary for problems like local broadcast. Thus we can observe concrete tradeoffs depending on model assumptions.

Theorem 5.3. *For every (possibly randomized) broadcast algorithm \mathcal{A} that uses neither node coordinates nor NTD primitive, there is a $(\epsilon R/8, 1)$ -bounded-independence metric space where \mathcal{A} needs $\Omega(n)$ rounds to do broadcast in a $O(1)$ -broadcastable network, even if the nodes have CD and ACK primitives and operate spontaneously.*

References

- [1] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. On the Complexity of Radio Communication. In *STOC*, pages 274–285, 1989.
- [2] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. A lower bound for radio broadcast. *Journal of Computer and System Sciences*, 43(2):290–298, 1991.
- [3] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in radio networks: an exponential gap between determinism and randomization. In *PODC’87*, 1987.
- [4] L. Barenboim and D. Peleg. Nearly optimal local broadcasting in the SINR model with feedback. In *SIROCCO*, pages 164–178, 2015.
- [5] M. H. Bodlaender and M. M. Halldórsson. Beyond geometry: towards fully realistic wireless models. In *PODC*, pages 347–356, 2014.
- [6] M. H. Bodlaender, M. M. Halldórsson, and P. Mitra. Connectivity and Aggregation in Multihop Wireless Networks. In *PODC*, pages 355–364. ACM, 2013.
- [7] I. Chlamtac and S. Kutten. On broadcasting in radio networks—problem analysis and protocol design. *IEEE Transactions on Communications*, 33(12):1240–1246, 1985.
- [8] B. S. Chlebus, D. R. Kowalski, and S. Vaya. Distributed communication in bare-bones wireless networks. *arXiv preprint arXiv:1510.07357*, 2015.
- [9] B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit disk graphs. *Discrete Math.*, 86(1):165–177, 1990.
- [10] A. E. Clementi, A. Monti, and R. Silvestri. Round robin is optimal for fault-tolerant broadcasting on wireless networks. *Journal of Parallel and Distributed Computing*, 64(1):89–96, 2004.
- [11] A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. In *FOCS’03*, 2003.
- [12] S. Daum, S. Gilbert, F. Kuhn, and C. Newport. Broadcast in the Ad Hoc SINR Model. In *DISC*, pages 358–372, 2013.
- [13] B. Derbel and E. Talbi. Radio Network Distributed Algorithms in the Unknown Neighborhood Model. In *ICDCN’10*, pages 155–166, 2010.
- [14] S. Dolev, S. Gilbert, M. Khabbazi, and C. Newport. Leveraging channel diversity to gain efficiency and robustness for wireless broadcast. In *DISC*, pages 252–267, 2011.
- [15] A. Fanghänel, T. Kesselheim, and B. Vöcking. Improved algorithms for latency minimization in wireless networks. *Theor. Comput. Sci.*, 412(24):2657–2667, 2011.
- [16] M. Ghaffari, S. Gilbert, C. Newport, and H. Tan. Optimal broadcast in shared spectrum radio networks. In *OPDIS’12*, 2012.
- [17] M. Ghaffari, B. Haeupler, and M. Khabbazi. Randomized broadcast in radio networks with collision detection. In *PODC’13*, 2013.

- [18] M. Ghaffari, B. Haeupler, N. Lynch, and C. Newport. Bounds on contention management in radio networks. In *DISC*, 2012.
- [19] M. Ghaffari, E. Kantor, N. Lynch, and C. Newport. Multi-message broadcast with abstract MAC layers and unreliable links. In *PODC'14*, 2014.
- [20] M. Ghaffari, N. Lynch, and C. Newport. The cost of radio network broadcast for different models of unreliable links. In *PODC'13*, 2013.
- [21] O. Goussevskaia, T. Moscibroda, and R. Wattenhofer. Local broadcasting in the physical interference model. In *DIALM-POMC '08*, pages 35–44, 2008.
- [22] P. Gupta and P. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.
- [23] M. M. Halldórsson. Wireless Scheduling with Power Control. *ACM Trans. Algorithms*, 9(1):7:1–7:20, Dec. 2012.
- [24] M. M. Halldórsson, S. Holzer, and N. A. Lynch. A local broadcast layer for the SINR network model. In *PODC*, pages 129–138, 2015.
- [25] M. M. Halldórsson, S. Holzer, P. Mitra, and R. Wattenhofer. The Power of Non-Uniform Wireless Power. In *SODA'13*, pages 1595–1606, 2013.
- [26] M. M. Halldórsson and P. Mitra. Nearly Optimal Bounds for Distributed Wireless Scheduling in the SINR Model. In *ICALP*, pages 625–636, 2011.
- [27] M. M. Halldórsson and P. Mitra. Distributed Connectivity of Wireless Networks. In *PODC*, 2012.
- [28] M. M. Halldórsson and P. Mitra. Towards tight bounds for local broadcasting. In *FOMC*, 2012.
- [29] M. M. Halldórsson, Y. Wang, and D. Yu. Leveraging multiple channels in ad hoc networks. In *PODC'15*, 2015.
- [30] M. Hoefer, T. Kesselheim, and B. Vöcking. Approximation algorithms for secondary spectrum auctions. In *Proceedings of the Twenty-third Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '11, pages 177–186, New York, NY, USA, 2011. ACM.
- [31] R. Impagliazzo and V. Kabanets. Constructive proofs of concentration bounds. In *APPROX-RANDOM*, pages 617–631, 2010.
- [32] T. Jurdzinski, D. Kowalski, M. Rozanski, and G. Stachowiak. Distributed randomized broadcasting in wireless networks under the SINR model. In *DISC*, pages 373–387, 2013.
- [33] T. Jurdzinski, D. Kowalski, M. Rozanski, and G. Stachowiak. On the impact of geometry on ad hoc communication in wireless networks. In *PODC '14*, 2014.
- [34] T. Jurdzinski, D. R. Kowalski, and G. Stachowiak. Distributed Deterministic Broadcasting in Uniform-power Ad Hoc Wireless Networks. In *FCT'13*, pages 195–209, 2013.
- [35] T. Jurdzinski, D. R. Kowalski, and G. Stachowiak. Distributed deterministic broadcasting in wireless networks of weak devices. In *ICALP'13*, 2013.

- [36] T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *DISC*, pages 163–178, 2010.
- [37] D. Kowalski and A. Pelc. Broadcasting in undirected ad hoc radio networks. In *PODC'03*, 2003.
- [38] F. Kuhn, N. Lynch, and C. Newport. Brief announcement: Hardness of broadcasting in wireless networks with unreliable communication. In *DISC*, 2009.
- [39] F. Kuhn, N. Lynch, C. Newport, R. Oshman, and A. Richa. Broadcasting in Unreliable Radio Networks. In *PODC*, pages 336–345. ACM, 2010.
- [40] F. Kuhn, T. Moscibroda, and R. Wattenhofer. Initializing Newly Deployed Ad Hoc and Sensor Networks. In *MobiCom*, pages 260–274. ACM, 2004.
- [41] F. Kuhn and R. Oshman. Dynamic networks: models and algorithms. *ACM SIGACT News*, 42(1):82–96, 2011.
- [42] F. Kuhn, R. Wattenhofer, and A. Zollinger. Ad hoc networks beyond unit disk graphs. *Wireless Networks*, 14:715–729, 2008.
- [43] E. Kushilevitz and Y. Mansour. An $\Omega(D \log(n/D))$ lower bound for broadcast in radio networks. *SIAM J. Comput.*, 27(3):702–712, 1998.
- [44] C. Lenzen and R. Wattenhofer. Distributed algorithms for sensor networks. *Philosophical Transactions of the Royal Society A*, 370(1958), 2012.
- [45] T. Moscibroda and R. Wattenhofer. Coloring unstructured radio networks. *Distributed Computing*, 21(4):271–284, 2008.
- [46] C. Newport. Radio network lower bounds made easy. In *DISC'14*, 2014.
- [47] C. Scheideler, A. Richa, and P. Santi. An $O(\log N)$ Dominating Set Protocol for Wireless Ad-hoc Networks Under the Physical Interference Model. In *MobiHoc '08*, pages 91–100, 2008.
- [48] S. Schmid and R. Wattenhofer. Algorithmic models for sensor networks. In *WPDRTS'06*, 2006.
- [49] J. Schneider and R. Wattenhofer. Coloring Unstructured Wireless Multi-hop Networks. In *PODC*, pages 210–219, 2009.
- [50] J. Schneider and R. Wattenhofer. What is the Use of Collision Detection (in Wireless Networks)? In *DISC*, pages 133–147. Springer-Verlag, 2010.
- [51] D. Yu, Q.-S. Hua, Y. Wang, and F. Lau. An $O(\log n)$ Distributed Approximation Algorithm for Local Broadcasting in Unstructured Wireless Networks. In *DCOSS '12*, pages 132–139, 2012.
- [52] D. Yu, Q.-S. Hua, Y. Wang, H. Tan, and F. Lau. Distributed multiple-message broadcast in wireless ad-hoc networks under the SINR model. In *SIROCCO'12*, pages 111–122, 2012.
- [53] D. Yu, Q.-S. Hua, Y. Wang, J. Yu, and F. Lau. Efficient distributed multiple-message broadcasting in unstructured wireless networks. In *INFOCOM'13*, pages 2427–2435, 2013.

- [54] D. Yu, Y. Wang, Q.-S. Hua, and F. C. M. Lau. Distributed local broadcasting algorithms in the physical interference model. In *Distributed Computing in Sensor Systems and Workshops (DCOSS), 2011 International Conference on*, pages 1–8, 2011.
- [55] D. Yu, Y. Wang, Y. Yan, J. Yu, and F. Lau. Speedup of information exchange using multiple channels in wireless ad hoc networks. In *INFOCOM'15*, 2015.

A Other Related Work

Wireless models: Considering wireless interference, there are two classes of wireless network models: graph-based and physical models. Basically, the graph-based models define a local and binary type of interference, while the physical models consider fading effect of signal on wireless channels. The most classical graph-based model is the radio network model [7]. In this model, the network is modeled using a communication graph, where each pair of nodes that can communicate with each other is connected by an edge. It defines the interference just from direct neighbors, and a transmission can succeed if and only if there is only one neighbor of the receiver transmitting. There are many widely used variants of the classical radio network models, including: 1) the k -hop model where the interference comes from k -hop neighbors [48]; 2) Unit Disc Graph (UDG) model [9] which defines the neighborhood using a unit disc; 3) Quasi Unit Disc Graph (QUDG) model [42] which just defines all pairs of nodes with distance at most ρ for some given $\rho \in (0, 1]$ are adjacent, and leave the ‘grey’ area in $(\rho, 1]$ being determined by an adversary; 4) Protocol model [22], where each node has a transmission range and an interference range, and a successful transmission occurs if a node falls into the transmission range of a transmitter and outside the interference ranges of all other transmitter; 5) Bounded-Independence Graph (BIG) model [50], which defines abstractly and requires that the size of the maximal independent set in the r -hop neighborhood of each node is bounded by a polynomial function with r . Though the graph-based interference models miss certain crucial aspects of actual wireless networks, the simple definition of these models can help derive novel insights into distributed solutions to wireless problems.

Physical models, also known as SINR models [22], capture the fading and cumulative features of receptions in actual wireless environments. The default assumption is that interference fades with a polynomial of the distance, and transmission succeeds only if the received signal strength is sufficiently larger than the total interference plus noise. Recently, the SINR model has attracted great attentions in the distributed community [6, 5, 12, 21, 25, 26, 27, 29, 33, 34, 36, 47, 51, 53, 54, 55]. Most of these works focus on networks embedded in Euclidean space, while many of the results hold also for doubling or “bounded growth” metrics [5, 12, 33]. Those doubling metrics constrain growth at every (or arbitrarily small) granularity, while ours only bounds regions proportional to the transmission range, in particular capturing bounded independence graphs (BIG). For more on wireless models, please refer to [44, 48].

Local Broadcast: In the radio network model, probably the first local broadcast result was a randomized algorithm of Alon et al. [1] in a synchronous model, running in $O(\Delta \log n)$ rounds. Derbel and Talbi [13] later generalized their algorithm to work without knowledge of Δ and their proposed algorithm can accomplish local broadcast in $O(\Delta \log n + \log^2 n)$ rounds. The decay strategy also yields an $O(\Delta \log n)$ time algorithm for local broadcast without knowledge of Δ [3, 18].

Goussevskaia et al. [21] gave the first results for local broadcast in the SINR model, running in time $O(\Delta \log n)$ and $O(\Delta \log^3 n)$ with and without knowledge of Δ , respectively. The latter was improved in [54] and further improved, independently, to $O(\Delta \log n + \log^2 n)$ time [28, 51]. With free acknowledgements, this was improved to $O(\Delta + \log^2 n)$ [28]. When additionally Δ is known,

this was further improved recently to $O(\Delta + \log n \cdot \log \log n)$ in the spontaneous setting [4]. The speedup of multiple channels on local broadcast was considered in [29, 55]

Broadcast: The complexity of broadcasting is well understood in graph-based models. In the radio network model, Bar-Yehuda et al. [3] presented the decay protocol which can accomplish non-spontaneous broadcast in $O(D \log n + \log^2 n)$ rounds, where D is the diameter. This result was improved to $\Theta(D + \log n) \log(n/D)$ independently by Czumaj and Rytter [11], and Kowalski and Pelc [37]. These algorithms can be viewed as clever optimizations of the decay protocol and match the lower bound [2, 43, 46]. With collision detection, this lower bound was recently broken in [17], where a solution of $O(D + \text{polylog}(n))$ was given. Broadcast in multi-channel radio networks was considered in [14, 16]. For the UDG model, an $O(D + \log^2 n)$ time algorithm was given in [12] in the spontaneous setting.

Distributed models of temporal variability: Dynamic networks have been studied extensively in recent years (see [41] for a survey), but generally not in the presence of interference.

The *dual graph* model [38] (originally due to [10]) was designed to capture inherent *unreliability* in wireless networks, much of which can be due to dynamicity. The main focus of that work is on extending the radio network model in general graphs. Importantly, the dual graph model does not distinguish between interference and communication edges; only that the unreliable edges can transmit both interference and the usual communication, but their availability is under adversarial control. Thus, there is no way to capture interference from further away nodes. Most problems become extremely difficult against a powerful adversary, and to get good result, one must assume a much weaker one [20]. This model only involves edge behavior and not node changes (churn). Both the local broadcast [18] and global broadcast [38, 39, 20, 19] problems are studied in the dual graph model.

A dynamic model that considers node insertion is the *unstructured* model [40], which admits arbitrary wake-up mode and asynchronous communication. This model was first proposed in the unit-disc setting, and then extended to bounded independence graphs (BIG) [49] and SINR [21]. It has been widely used in the solution of a variety of distributed wireless problems [40, 45, 49, 50, 28, 21, 54, 51, 55], including local broadcast [28, 21, 54, 51, 55], but there are no known global broadcast results. The model neither consider node deletion nor edge changes. Hence, the impact of dynamicity on global communication is largely unexplored.

B Implementing Communication and Primitives

Modeling communication Our communication model captures most known algorithmic wireless models, as it is demonstrated below on the example of SINR, disk-graph based and Protocol models.

Note that we assume below the distance $d(x, y)$ to be symmetric, but the results hold also for “almost symmetric” functions, i.e. when there is a constant c such that $d(x, y) \leq c \cdot d(y, x)$ for all x, y .

SINR Model. Consider a network in a metric space. In the SINR model of communication, if S is the set of simultaneously transmitting nodes in the network, a node v receives the transmission of node u if and only if

$$\frac{P}{d(u, v)^\zeta} > \beta \cdot \left(\sum_{w \in S \setminus \{u, v\}} \mathcal{I}_{wv} + N \right),$$

where constants $\beta \geq 1$ and $N > 0$ denote the minimum SINR threshold and the ambient noise, respectively.

Note that $R = (P/(\beta N))^{1/\zeta}$ in this setting. We can implement SuccClear with parameters $\mathcal{I}_c = \min\{\beta, ((1 - \epsilon)^{-\zeta} - 1)\}N/2^\zeta$ and $\rho_c = 0$, as shown in the proposition below.

Proposition B.1. *If the interference at a node v is less than \mathcal{I}_c in round t , then v will deliver its message if it transmits.*

Proof. Note that if the interference at node v is not more than \mathcal{I}_c then there is no node $u \in D_v^2$ transmitting in round t , as otherwise the interference at v would be at least $P/(2R)^\zeta = \beta N/2^\zeta$; hence, only nodes in \bar{D}_v^2 can transmit. Consider an arbitrary node $w \in N(v)$. For each node $u \in \bar{D}_v^2$, $d(u, w) \geq d(u, v) - d(v, w) \geq d(u, v)/2$. Then, the interference at w is at most:

$$\mathcal{I}_w \leq \sum_{u \in \bar{D}_v^2} \frac{P}{d(u, w)^\zeta} \leq 2^\zeta \sum_{u \in \bar{D}_v^2} \frac{P}{d(u, v)^\zeta} \leq 2^\zeta \mathcal{I}_c \leq ((1 - \epsilon)^{-\zeta} - 1)N,$$

which implies that node w receives v 's transmission: $\frac{P/d(v, w)^\zeta}{N + \mathcal{I}_w} \geq \frac{\beta N(1 - \epsilon)^{-\zeta}}{N + ((1 - \epsilon)^{-\zeta} - 1)N} \geq \beta$. \square

The UDG and UBG Models. These models are described by geometric graphs: a node u receives a message from another node v if and only if v is the only transmitting neighbor of u .

In the Unit Disk Graph (UDG) and Unit Ball Graph (UBG) models the nodes are located in a metric space and two nodes are connected by an edge if and only if their distance is at most R . The functionality of SuccClear can be modeled as follows: the transmission of a node v is received by all its neighbors if there is no other node at distance less than $2R$ from v transmitting simultaneously, i.e. we can set the parameters to $\mathcal{I}_c = \infty$ and $\rho_c = 2$.

The Quasi-UDG Model. The Quasi-UDG model is an extension of the UDG model: a) if $d(u, v) \leq R$ then u and v are connected by an edge, b) if $d(u, v) > R'$ then they are disconnected, c) otherwise, u and v may be connected or not. In this case SuccClear may be implemented by setting $\mathcal{I}_c = \infty$ and $\rho_c = (R + R')/R$, with the adversary constrained to follow the specific static situation captured by the QUDG.

The Protocol Model. In the Protocol Model, the nodes are in a metric space and there are two radii: R – the communication radius, and R' – the interference radius. A node v receives the transmission of a node u if and only if: 1) v is in the communication range of node u , i.e. $d(u, v) \leq R$, and 2) there is no transmitting node w such that v is in the interference range of w : for each transmitting node $w \neq u$, $d(w, v) > R'$. SuccClear may be implemented here by setting $\mathcal{I}_c = \infty$ and $\rho_c = (R + R')/R$.

The BIG Model. In the Bounded Independence Graph (BIG) model, for a parameter λ , we are given a graph on the nodes with the property that for every node v and every $k \geq 1$, the maximum independent set in the k -neighborhood of v is $O(k^\lambda)$. The shortest-path distance metric on the graph is now naturally a $(1, \lambda)$ -bounded independence metric. To fit in our model, the growth parameter λ must be less than ζ .

k-hop Variants. These graph models can be naturally generalized to a model on interference, where nodes of distance at most k cause interference, for some $k > 1$. We capture this by extending ρ_c as needed.

Implementing primitives with physical carrier sensing One way of implementing the primitives mentioned in this paper is to use physical carrier sensing, i.e. we assume the nodes have technology to detect if the interference (plus noise) is higher than a given threshold.

We show below how to implement ACK, CD and NTD primitives using carrier sensing.

CD primitive. The CD primitive can be implemented using a carrier sensing threshold $T = P/((1 - \epsilon)R)^\zeta$ and setting the parameter $\mathcal{I}_{cd} < T$; BUSY channel is detected if and only if the interference is at least T .

We will need the following technical fact.

Lemma B.2. *For every $x_i \in [0, \frac{1}{2}]$, $i = 1, 2, \dots, n$, it holds that $4^{-\sum_{i=1}^n x_i} \leq \prod_{i=1}^n (1 - x_i) \leq e^{-\sum_{i=1}^n x_i}$.*

Proposition B.3. *If $P_t(v) > \phi \geq 1$ in round t then all nodes in $B(v, R/2)$ detect BUSY channel with probability at least $1 - (1 + 2\phi)e^{-\phi}$. In particular, we can take $h_1 = 2$ if $\phi \geq 10$.*

Proof. By the setting of T and the definition of $B(v, R/2)$, if two nodes in $B(v, R/2)$ transmit in round t , then all nodes in $B(v, R/2)$ will detect BUSY channel. Hence, the probability of all nodes in $B(v, R/2)$ detecting BUSY channel is at least the probability of more than one node transmitting in round t . The probability of no node transmitting is $p_0 = \prod_{u \in B(v, R/2)} (1 - p_t(u)) \leq e^{-\phi}$ by Lemma B.2. The probability of exactly one node transmitting is:

$$p_1 = \sum_{u \in B(v, R/2)} p_t(u) \prod_{w \in B(v, R/2) \setminus u} (1 - p_t(w)) \leq 2 \sum_{u \in B(v, R/2)} p_t(u) \prod_{w \in B(v, R/2)} (1 - p_t(w)) \leq 2\phi e^{-\phi},$$

where we used the assumption that $p_t(u) \leq 1/2$ and Lemma B.2. Thus, the probability of detecting BUSY channel is at least $1 - p_0 - p_1 \geq 1 - (1 + 2\phi)e^{-\phi}$. \square

Proposition B.4. *For every $\rho > 0$, if $P_t^\rho(v) < \eta$ and $\mathcal{I}_t^\rho(v) < \mathcal{I}_{cd}$ then v detects IDLE channel with probability at least $4^{-\eta}$.*

Proof. By the setting of the threshold T , if there is no node transmitting in D_v^ρ then node v will detect IDLE channel. Thus, the probability that v detects IDLE channel is at least

$$\prod_{u \in D_v^\rho} (1 - p_t(u)) \geq 4^{-\sum_{u \in D_v^\rho} p_t(u)} \geq 4^{-\eta},$$

using Lemma B.2. \square

ACK primitive. In order to detect successful transmission, we can use interference threshold $T = \min\{\mathcal{I}_c, P/(\rho_c R)^\zeta\}$ and set $\mathcal{I}_{ack} < T$, where $\mathcal{I}_c(\epsilon)$ and $\rho_c(\epsilon)$ are the parameters of SuccClear. If a node v senses that the interference is no higher than T , it knows that: 1. there is no node in $D_v^{\rho_c}$ transmitting, as otherwise the interference would be at least $P/(\rho_c R)^\zeta$, 2. the interference is at most \mathcal{I}_c ; thus, it knows that its transmission has been received by all neighbors in $N_t(\epsilon)$ by SuccClear.

NTD primitive. If a node v receives a message from a node u then v can separate the signal from the interference and measure the received signal strength. As the nodes use uniform power assignment, node v knows that $u \in D_v^{\epsilon/2}$ if the received signal is stronger than $P/(\epsilon R/2)^\zeta$.

Implementing primitives by other means The primitives can frequently be implemented in other ways, often with the logarithmic blowup that explains the differences with the best carrier-sense-free results.

CD primitive. In an asynchronous system, it may be impossible to implement CD by other means than carrier sense. In a synchronized system, however, we can be achieved with logarithmic or polylogarithmic factor overhead. Consider a given round. For each probability $p = 2^{-i}$, $i = 1, 2, \dots, \log n$, repeat $C \log n$ times: the senders in the original round transmit with probability p .

Using concentration bound with C sufficiently large, one can infer the contention within an small approximation, with high probability. Such a strategy has been applied, e.g., in [27].

ACK primitive. A simple strategy is to work with only probabilistic guarantees of a transmission being received by all neighbors. Then, simply repeat the protocol until this has been achieved $C \log n$ times, which gives an ACK guarantee, w.h.p. This approach underlies, e.g., the local broadcast algorithms without carrier sense [21, 28, 51].

NTD primitive. This primitive, which is essential for dominator-based strategies for broadcast, can be implemented using power control: by lowering the power on all units appropriately, one can ensure that nodes further away (by a small constant factor) will not be able to hear the message due to the ambient noise term, see e.g. [52]. Alternatively, one can assume that distances can be determined in other ways, such as by GPS [32, 34].

C Proof of Proposition 3.1: Contention Control

Recall that we need to prove the following. We assume that at the beginning of the *first* phase under consideration, the contention in the whole network is bounded by a constant. This holds for all algorithms in this paper, as the initial probability of nodes is always at most $1/n$.

Proposition 3.1. Let $\sigma \in (0, 1)$. If constants $\rho = \rho(\sigma, \hat{\mathcal{I}})$, $\hat{\eta} = \hat{\eta}(\rho, \sigma)$ and $\gamma = \gamma(\rho, \hat{\eta}, \sigma, \hat{\mathcal{I}})$ are large enough then for each node v and phase H , with probability $1 - O(n^{-3})$, a $(1 - \sigma)$ -fraction of the rounds in H are good.

Recall that in a good round, there should be both bounded contention and low interference. The proof is split into two parts, each handling one of these properties. Prop. 3.1 follows by simply combining those two parts.

We will need the following concentration bounds.

Lemma C.1. [47, 31] Consider a collection of binary random variables X_1, \dots, X_n , and let $X = \sum_{i=1}^n X_i$. If there are probabilities p_1, \dots, p_n with $E[\prod_{i \in S} X_i] \leq \prod_{i \in S} p_i$ for every set $S \subseteq \{1, \dots, n\}$, then it holds for $\mu = \sum_{i=1}^n p_i$ and $\delta > 0$ that

$$\Pr[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu \leq e^{-\frac{\delta^2 \mu}{2(1+\delta/3)}}.$$

If, on the other side, there are probabilities p_1, \dots, p_n with $E[\prod_{i \in S} X_i] \geq \prod_{i \in S} p_i$ for every set $S \subseteq \{1, \dots, n\}$, then it holds for $\mu = \sum_{i=1}^n p_i$ and $0 < \delta < 1$ that

$$\Pr[X \leq (1 - \delta)\mu] \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^\mu \leq e^{-\delta^2 \mu / 2}.$$

C.1 Bounded Contention Rounds

First, we show that for each fixed node u , the contention in the *local* neighborhood $B(u, R/2)$ is bounded in most of the rounds of a phase. To this end, we show that in each round, the contention is either already low or will be halved with significant probability, then apply a concentration bound to show the claim. Recall the constant h_1 from the definition of CD.

Lemma C.2. Let H be a time interval and assume that $P_t(u) = \phi_0 \geq 0$ at the beginning of H . Then for every $\phi \geq 3$, $P_t(u) > \phi$ happens at most $O(\phi^{-k}) \cdot |H| + \max\{0, \log_{6/5}(\phi_0/\phi)\}$ times during H , with probability $1 - 2^{-\mu|H|}$, where $\mu = h_1^{-\phi} / (1 - O(\phi^{-k}))$ and $k > 0$ is the edge change parameter.

Proof. By the assumptions on edge changes, we have that during phase H , the fraction of rounds where more than $\phi/4$ nodes become a neighbor of u because of edge changes is $O(\phi^{-k})$. Let H' denote the remaining set of rounds. We have $|H'| = (1 - O(\phi^{-k}))|H|$. In each of those rounds, the contribution of edge changes in $P_t(u)$ is clearly at most $\phi/8$.

Next we bound the number of rounds in H' , where $P_t(u) > \phi$. Consider such a round $t \in H'$. Let $\mu' = h_1^{-\phi}$. By the definition of CD, all nodes in $B(u, R/2)$ detect BUSY channel in round t and halve their transmission probabilities with probability at least $1 - \mu'$. Thus, $Pr[P_{t+1}(u) \leq \frac{1}{2}P_t(u) + 1/2 + \phi/8] \geq 1 - \mu'$, where the additive $1/2$ accounts for the sum of probabilities of the nodes that just join the ball $B(u, R/2)$ due to node churn (recall that each of them has an initial probability at most $1/2n$) and $\phi/8$ is an upper bound on the contention due to edge changes (by the definition of H'). For each round $t \in H'$, let us define a binary random variable X_t as follows: $X_t = 0$ if $P_t(u) \leq \phi$ or $P_t(u) > \phi$ and $P_{t+1}(u) \leq \frac{1}{2}P_t(u) + 1/2 + \phi/8 \leq 5P_t(u)/6$ (recall that $\phi \geq 3$) and $X_t = 1$ otherwise. By the discussion above, we have:

$$\begin{aligned} Pr[X_t = 0] &= Pr[P_t(u) \leq \phi] + Pr[P_t(u) > \phi \wedge P_{t+1}(u) \leq \frac{5}{6}P_t(u)] \\ &= Pr[P_t(u) \leq \phi] + Pr[P_t(u) > \phi] \cdot Pr[P_{t+1}(u) \leq \frac{5}{6}P_t(u) | P_t(u) > \phi] \\ &\geq Pr[P_t(u) \leq \phi] + Pr[P_t(u) > \phi] \cdot (1 - \mu') \\ &\geq 1 - \mu'. \end{aligned}$$

This implies that for each round $t \in H$ and each subset $S \subseteq H$ of earlier rounds ($s < t$ for $s \in S$), $Pr[X_t = 0 | \prod_{s \in S} X_s = 1] \geq 1 - \mu'$ and $Pr[X_t = 1 | \prod_{s \in S} X_s = 1] \leq \mu'$. Thus, for each subset $S \subseteq H'$ we get $E[\prod_{s \in S} X_s] \leq \mu'^{|S|}$, and if we denote $X = \sum_{t \in H'} X_t$, then $E[X] \leq \mu'|H'| = \mu|H|$, where we denote $\mu = \mu'|H|/|H'|$. Thus, we can apply Chernoff bound (Lemma C.1) with $\delta = 3/2$ to bound X with high probability:

$$Pr[X \geq 3\mu|H|/2] \leq 2^{-\mu|H|}. \quad (1)$$

Using this bound, we obtain a bound on the number of rounds t with $P_t(u) > \phi$. Let Y denote this number. Consider a maximal interval $\hat{H} = [r_1, r_2] \subseteq H'$ such that $P_t(u) > \phi$ for all $t \in \hat{H}$. For each round $t \in \hat{H}$, $P_{t+1}(u) \leq \frac{5}{6}P_t(u)$ if $X_t = 0$ and $P_{t+1}(u) \leq 2P_t(u) + 1 + \phi/4 < 3P_t(u)$ otherwise. By maximality of \hat{H} , if r_1 is not the first round of phase H , then we have $P_{r_1-1}(u) \leq \phi$. Then a simple calculation shows that $X_t = 1$ for at least $1/8$ part of the rounds of \hat{H} , i.e. $\hat{X} = \sum_{t \in \hat{H}} X_t \geq |\hat{H}|/8$. On the other hand, if there is a unique maximal interval \hat{H} starting at the first round of phase H and such that $P_{r_1}(u) > \phi$, then it must hold that $\hat{X} \geq (|\hat{H}| - L)/8$, where $L = \max\{0, \log_{6/5}(\phi_0/\phi)\}$. Combining these observations, we get that $Y \leq 8X + L$.

By combining this bound with (1), we have:

$$Pr[Y \geq 8 \cdot 3\mu|H|/2 + L] \leq Pr[X \geq \delta\mu|H|] < 2^{-\mu|H|}.$$

Thus, with probability $1 - 2^{-\mu|H|}$, in at most $O(h_1^{-\phi} + \phi^{-k}) = O(\phi^{-k})$ fraction of rounds in H there can be $P_t(u) > \phi$. \square

Recall that for node v we need to show that $P_t^\rho(v) \leq \hat{\eta}$ for most rounds in a phase. We prove this by taking a *constant size* $R/2$ -cover of D_v^ρ and applying Lemma C.2 to all nodes in the cover *simultaneously*.

Proposition C.3. *For every $\sigma \in (0, 1)$ and $\rho > 0$, there are constants $\hat{\eta} = \hat{\eta}(\rho, \sigma)$ and $\gamma = \gamma(\rho, \sigma)$, such that for every node v and phase H , with probability $1 - O(n^{-3})$, at least $(1 - \sigma)$ -fraction of rounds in H are bounded contention rounds, i.e. $P_t^\rho(v) \leq \hat{\eta}$.*

Proof. It will be sufficient to choose γ and $\hat{\eta}$ such that $\hat{\eta} = O(\rho^{\lambda+\lambda/k}/\sigma^{1/k})$ and $\gamma = h_1^{O(\hat{\eta}/\rho^\lambda)}$. By the definition of λ and bounded independence, there is a $R/2$ -cover S of D_v^ρ of size $|S| = O(\rho^\lambda)$. Set $\phi = \hat{\eta}/|S|$. Let us fix a node $u \in S$. If H is the first phase or $P_t(u) \leq \phi$ then we set $\hat{H} = H$. Otherwise, consider the phase H' of $\gamma \log n$ rounds preceding H . By Lemma C.2, if we set $\gamma \geq 3h_1^\phi$ then, with probability $1 - n^{-3}$, there is a round in H' where $P_t(u) \leq \phi$. Let t_0 be such a round and set $\hat{H} = [t_0, t_l]$, where t_l is the last round of H . Clearly, $|\hat{H}| \leq 2|H|$. Now we can apply Lemma C.2 for node u with ϕ as above and conclude that with probability at least $1 - n^{-3}$, there are at most $O(\phi^{-k}) \cdot |\hat{H}|$ rounds in \hat{H} (note that $\phi_0 = \phi$) where $P_t(u) > \phi$, so there are at most $O(\phi^{-k}) \cdot |H|$ rounds in H where $P_t(u) > \phi$. Now we can apply the same argument for all nodes in $|S|$ simultaneously and conclude that with probability at least $1 - O(n^{-3})$, there are at most $O(\phi^{-k}) \cdot |S| |H| \leq \sigma |H|$ rounds in H where $P_t(u) > \phi$ for every $u \in S$, if we take $\hat{\eta} = O(\rho^{\lambda+\lambda/k}/\sigma^{1/k})$ to be large enough. In the remaining $(1 - \sigma)|H|$ rounds we have $P_t^\rho(v) \leq \sum_{u \in S} P_t(u) \leq \hat{\eta}$. \square

We extract the following result from the proof of Lemma C.2, to use it later. The proof is similar to the proof of Prop. C.3.

Corollary C.4. *For every $\phi \geq 3$ and each node v , the expected number of rounds in each phase H where $P_t(v) > \phi$, is $O(\phi^{-k}) \cdot |H|$.*

Proof. Let \hat{H} be the shortest time interval containing H such that $P_t(v) \leq \phi$ at the beginning of \hat{H} . Let \mathcal{E} denote the event that $|\hat{H}| \leq 2|H|$. As in the proof of Prop. C.3, we have $\Pr[\mathcal{E}] > 1 - n^{-3}$. Let X denote the number of rounds $t \in \hat{H}$ with $P_t(v) > \phi$. We know from the proof of Lemma C.2 that $E[X|\mathcal{E}] = O(\phi^{-k}) \cdot |\hat{H}| = O(\phi^{-k}) \cdot |H|$ and $E[X] = E[X|\mathcal{E}] \cdot \Pr[\mathcal{E}] + E[X|\bar{\mathcal{E}}] \cdot \Pr[\bar{\mathcal{E}}] = O(\phi^{-k}) \cdot |H|$. This completes the proof as $H \subseteq \hat{H}$. \square

C.2 Low Interference Rounds

Next we show that if ρ is appropriately chosen then for each node v , $\hat{\mathcal{I}}_t^\rho(v) < \hat{\mathcal{I}}$ happens most of the time during each phase w.h.p. In order to show this, first we split the set of nodes in \bar{D}_v^ρ into local neighborhoods. We show that if in a given round the contention in each local neighborhood is bounded by an appropriate threshold, then the expected interference is small. For each local neighborhood, the expected number of rounds when the contention is higher than its threshold is bounded using Corollary C.4. Combining these results into one we get a bound on the expected number of rounds when $\hat{\mathcal{I}}_t^\rho(v)$ is small. It then remains to apply a concentration bound to conclude the proof.

Proposition C.5. *For every $\sigma \in (0, 1)$ and $\hat{\mathcal{I}} > 0$, there are constants $\rho = \rho(\sigma, \hat{\mathcal{I}})$ and $\gamma = \gamma(\sigma)$, such that for each node v and phase H , with probability $1 - O(n^{-3})$, at least $(1 - \sigma)$ -fraction of rounds in H are low interference rounds, i.e. $\hat{\mathcal{I}}_t^\rho(v) \leq \hat{\mathcal{I}}$.*

Proof. Let D_i denote the in-ball of radius $\rho R + iR$ centered at node v for $i = 0, 1, \dots$. Let $S_i \subseteq D_i$ be an $R/2$ -cover of D_i of size $|S_i| \leq C \cdot (\rho + i)^\lambda$ for a constant $C = C(\lambda)$, which exists by bounded independence. Let $\phi_i = (\rho + i)^{(\zeta-\lambda)/2} \phi$ for $i = 1, 2, \dots$, where $\phi = \Omega(\lambda/\sigma)$ is a large enough constant to be defined later.

Claim C.6. *If $P_t(u) < \phi_i$ holds for all i and $u \in S_i \setminus S_{i-1}$ and $\rho \geq \left(\frac{2C\phi P(\zeta+\lambda)}{R^\zeta \hat{\mathcal{I}}(\zeta-\lambda)} \right)^{2/(\zeta-\lambda)}$ then $\hat{\mathcal{I}}_t^\rho(v) \leq \hat{\mathcal{I}}$.*

Proof. We have, by the definition of expected interference (explanations below),

$$\begin{aligned}
\hat{\mathcal{I}}_t^\rho(v) &= \sum_{w \in \bar{D}_v^\rho} \frac{P}{d(w, v)^\zeta} \cdot p_t(w) \leq \sum_{i \geq 1} \frac{P}{((\rho + i - 1)R)^\zeta} \cdot |S_i \setminus S_{i-1}| \cdot \phi_i \\
&\leq \frac{\phi P(1 + 1/\rho)}{R^\zeta} \sum_{i \geq 1} \frac{|S_i \setminus S_{i-1}|}{(\rho + i)^{(\zeta + \lambda)/2}} \\
&\leq \frac{2\phi P}{R^\zeta} \sum_{i \geq 1} |S_i| \cdot \left(\frac{1}{(\rho + i)^{(\zeta + \lambda)/2}} - \frac{1}{(\rho + i + 1)^{(\zeta + \lambda)/2}} \right) \\
&\leq \frac{2\phi P}{R^\zeta} \sum_{i \geq 1} C(\rho + i)^\lambda \cdot \frac{(\zeta + \lambda)/2}{(\rho + i)^{(\zeta + \lambda)/2 + 1}} \\
&= \frac{C\phi P(\zeta + \lambda)}{R^\zeta} \sum_{i \geq 1} (\rho + i)^{-(1 + (\zeta - \lambda)/2)} \\
&\leq C\phi(\zeta + \lambda) \cdot (P/R^\zeta) \cdot 2/(\zeta - \lambda) \cdot \rho^{-(\zeta - \lambda)/2},
\end{aligned}$$

where the first inequality follows from the definition of sets S_i and the assumption of the claim, the second one follows by $\phi_i = (\rho + i)^{(\zeta - \lambda)/2} \phi$ and $(\rho + i - 1)(1 + 1/\rho) \geq \rho + i$, the third one is a rearrangement of the sum, the fourth one follows by the fact that $\frac{1}{(t-1)^\alpha} - \frac{1}{t^\alpha} < \frac{\alpha}{(t-1)^{\alpha+1}}$ for every $t > 1$ and $\alpha \geq 1$ (here, $(\zeta + \lambda)/2 > 1$), and the last one follows because $\zeta - \lambda > 0$. Thus, if $\rho \geq \left(\frac{2C\phi P(\zeta + \lambda)}{R^\zeta \hat{\mathcal{I}}(\zeta - \lambda)} \right)^{2/(\zeta - \lambda)}$ then the claim follows. \square

Let us put the nodes in an arbitrary order and let $Z_{i,j}$ denote the number of rounds $t \in H$ s.t. $P_t(u_j) \geq \phi_i$, where u_j is the j -th node in $S_i \setminus S_j$. By Corollary C.4, $E[Z_{i,j}] = O(\phi_i^{-k}) \cdot |H|$ holds for each i, j . Let Z denote the number of rounds in H when $\hat{\mathcal{I}}_t^\rho(v) > \hat{\mathcal{I}}$. By Claim C.6, $E[Z]$ can be bounded as follows,

$$E[Z] \leq \sum_i \sum_j E[Z_{i,j}] \leq \sum_i |S_i \setminus S_{i-1}| \cdot O(\phi_i^{-k}) \cdot |H| \leq \frac{\sigma}{3} |H|,$$

where the last inequality holds if $\phi = C''/\sigma^{1/k}$ for a large enough constant C'' ; we have $|S_i \setminus S_{i-1}| \leq C\lambda(\rho + i)^{\lambda-1}$ and $\phi_i^{-k} \leq \phi^{-k} \cdot (\rho + i)^{k(\zeta - \lambda)/2}$ with $k(\zeta - \lambda)/2 > \lambda$ (by definition), so we have $\sum_i |S_i \setminus S_{i-1}| \cdot O(\phi_i^{-k}) = O(\phi^{-k})$.

Each $Z_{i,j}$ can be considered as the sum of binary random variables similar to that in the proof of Lemma C.2. Thus, Z can also be seen as the sum of binary random variables. Moreover, the binary random variables satisfy the conditions of the Chernoff bound in Lemma C.1. Thus, by taking $\delta = 2$ and $\gamma \geq 9/\sigma$, we have: $\Pr[Z \geq 3 \cdot \frac{\sigma}{3} |H|] \leq 2^{-\sigma|H|/3} \leq n^{-3}$. \square

D Proof of Proposition 3.2: Transmissions in Good Rounds

Lemma D.1. *Let $\rho \geq \rho_c + 2$, $\hat{\mathcal{I}} \leq (1 - 1/\rho)^\zeta \mathcal{I}_c/10$ and $\phi > 0$, where ρ_c, \mathcal{I}_c are the parameters of SuccClear. Let t be a round such that for a node u , $P_t^\rho(u) < \phi$ and $\hat{\mathcal{I}}_t^\rho(u) < \hat{\mathcal{I}}$.*

If a node in $B(u, R/2)$ transmits in round t then it mass-delivers, with probability at least $0.9 \cdot 4^{-\phi}$.

Proof. Let v be the transmitting node, which could possibly be u itself. The probability that no other node in D_u^ρ transmits is

$$p' = \prod_{w \in D_u^\rho \setminus \{v\}} (1 - p_t(w)) \geq 4^{-\sum_{w \in D_u^\rho \setminus \{v\}} p_t(w)} \geq 4^{-\phi},$$

where we used Lemma B.2 with the assumption that $p_t(w) \leq 1/2$ and the bounded contention assumption.

On the other hand, we have that $\Pr[\mathcal{I}_t^{\rho-1}(v) \leq \mathcal{I}_c] \geq 9/10$. Indeed, $\Pr[\mathcal{I}_t^\rho(u) \leq \mathcal{I}' = (1 - 1/\rho)^\zeta \mathcal{I}_c] \geq 9/10$ follows by the bounded interference assumption and Markov inequality, so it suffices to show that $\mathcal{I}_t^{\rho-1}(v) \leq (1 - 1/\rho)^{-\zeta} \mathcal{I}_t^\rho(u)$, which we have because for each node $x \in \bar{D}_u^\rho$, $d(x, v) \geq d(x, u) - d(u, v) > (1 - 1/\rho)d_{xu}$, implying that

$$\mathcal{I}_t^{\rho-1}(v) = \sum_{x \in \bar{D}_u^\rho} P/d(x, v)^\zeta < \sum_{x \in \bar{D}_u^\rho} (1 - 1/\rho)^{-\zeta} \cdot P/d_{xu}^\zeta = (1 - 1/\rho)^{-\zeta} \mathcal{I}_t^\rho(u).$$

Thus, we have that with probability $0.9 \cdot 4^{-\phi}$, v transmits, no other node in D_u^ρ transmits and the interference from \bar{D}_u^ρ at v is at most \mathcal{I}_c , so v delivers its message, by SuccClear. \square

Proposition 3.2 *Assume that constants $\hat{\eta}, \rho, \gamma$ are large enough. For each node v and phase H , if at least $1/10$ -fraction of the rounds of H are of high contention, then $\Omega(|H|)$ nodes in D_v^ρ mass-deliver, with probability $1 - O(n^{-3})$.*

Proof. Let $\sigma = 1/10$. Let us choose constants $\hat{\eta}, \rho', \gamma$ such that Prop. 3.1 holds for $\sigma' = \sigma/2$ and $\hat{\mathcal{I}}_1 \leq (1/2 - 1/\rho')^\zeta \mathcal{I}_c/20$ and such that $\rho' > 2(\rho_c + 2)$ is large enough, as stated below, where ρ_c, \mathcal{I}_c are the parameters of SuccClear. We will use the fact that in Prop. 3.1, $\hat{\eta}$ has the following dependency on ρ' : $\hat{\eta} = O(\rho'^{\lambda + \lambda/k})$, where k is the edge change parameter (see the proof of Prop. C.3). We set $\rho = \rho'/2$. Namely, if we choose $\hat{\eta}, \rho, \gamma$ large enough (satisfying the relation above), then we have that with probability $1 - O(n^{-3})$, there are at least $(1 - \sigma/2)|H|$ rounds $t \in H$ such that $P_t^{2\rho}(v) < \hat{\eta}$ and $\hat{\mathcal{I}}_t^{2\rho}(v) < \hat{\mathcal{I}}_1$. Let t be such a round.

Claim D.2. *For each node $u \in D_v^\rho$, it holds that $P_t^\rho(u) < \hat{\eta}$ and $\hat{\mathcal{I}}_t^\rho(u) < \hat{\mathcal{I}}_2 = (1 - 1/\rho)^\zeta \mathcal{I}_c$.*

Proof. Let us fix a node $u \in D_v^\rho$. First note that $D_u^\rho \subseteq D_v^{2\rho}$, implying that $P_t^\rho(u) \leq P_t^{\rho'}(u) < \hat{\eta}$. It remains to bound the expected interference at u by nodes in \bar{D}_u^ρ . This interference can be split into two parts: 1. the interference by nodes in $\bar{D}_v^{\rho'}$ and 2. the interference by nodes in $D_v^{\rho'} \setminus D_u^\rho$. The first part can be bounded by $2^\zeta \hat{\mathcal{I}}_1 = \hat{\mathcal{I}}_2/2$ using a computation as in the proof of Lemma D.1. The second part can be bounded as follows: the contention in the area $D_v^{\rho'} \setminus D_u^\rho$ is at most $\hat{\eta}$ and the distance of those nodes to node u is at least ρ , so the expected interference by nodes in $D_v^{\rho'} \setminus D_u^\rho$ is at most $\hat{\eta} \cdot \frac{P}{(\rho R)^\zeta} = O(\frac{\rho^{\lambda + \lambda/k}}{\rho^\zeta}) = O(\frac{1}{\rho^{(\zeta - \lambda)/2}})$, and is less than $\hat{\mathcal{I}}_2/2$ if ρ is large enough, as $\zeta - \lambda > 0$. Here we used the assumption that $k > 2\lambda/(\zeta - \lambda)$. \square

Recall that there are at least $\sigma|H|$ high contention rounds for node v , i.e. rounds where $P_t^\rho(v) > \eta$. Let t be such a round. By the bounded independence property, there is a ρR -cover S of D_v^ρ of constant size. By the pigeonhole principle, there is a node $u \in S$ such that $P_t^\rho(u) > \eta' = \eta/|S|$.

By summarizing the above said, we get that there are at least $(1 - \sigma/2)|H|$ rounds t s.t. for all nodes $u \in D_v^\rho$, $P_t^\rho(u) < \hat{\eta}$ and $\hat{\mathcal{I}}_t^\rho(u) < \hat{\mathcal{I}}_2$. On the other hand, there are at least $\sigma|H|$ rounds where there is a node in D_v^ρ that has $P_t^\rho(u) > \eta'$. Thus, there are at least $\sigma|H|/2$ rounds t with the following property: there is a node $u \in D_v^\rho$ such that $P_t^\rho(u) \in (\eta', \hat{\eta})$ and $\hat{\mathcal{I}}_t^\rho(u) < \hat{\mathcal{I}}_2$. Fix such round t and node u . By Lemma D.1, there is a node $u' \in B(u, R/2)$ that transmits (the probability of this is at least η') and delivers its message with probability $0.9\eta'4^{-\hat{\eta}}$. Thus, the expected number of nodes in D_v^ρ that deliver their message during H is $\Omega(4^{-\hat{\eta}}|H|)$. The proof now follows by applying a Chernoff bound with the assumption that the constant γ is large enough: $\gamma = \Omega(4^{\hat{\eta}})$. \square

E Proof of Proposition 3.3: Low Contention Rounds

Proposition 3.3 *Assume that $\hat{\eta}, \rho, \gamma$ are large enough. For each node v and phase H , if at least $9/10$ -fraction of the rounds of H are low contention rounds, then with probability $1 - O(n^{-3})$, in at least $3/5$ -fraction of the rounds of H , v will detect IDLE channel and have low contention and low interference.*

Proof. Let $\sigma = 9/10$. Assume that $\hat{\eta}, \rho, \gamma$ large enough, such that Prop. 3.1 holds for $\sigma' = \sigma/10$ and $\hat{\mathcal{I}} \leq \mathcal{I}_{cd}/10$, where \mathcal{I}_{cd} is the parameter of CD. Let \mathcal{E} denote the event that there are at least $(1 - \sigma/10)|H|$ good rounds for v in phase H . By Prop. 3.1, we have $\Pr[\mathcal{E}] = 1 - O(n^{-3}) \geq 19/20$, if γ is large enough. Given \mathcal{E} , there are at least $9\sigma|H|/10$ rounds t that are both good and low contention for v . Let H' denote the set of such rounds and let X_t be a binary random variable with value 1 if and only if v detects IDLE channel in round $t \in H'$. Let $X = \sum_{t \in H'} X_t$. For each round $t \in H'$, since t is a good round and $\hat{\mathcal{I}} < \mathcal{I}_{cd}/10$, it holds with probability at least $9/10$ that the interference at node v from nodes in D_v^ρ is at most \mathcal{I}_{cd} . Given this and the fact that in low contention rounds we have $P_t^\rho(v) < \eta = \log_{h_2}(10/9)$, v will detect IDLE channel with probability at least $h_2^{-\eta} = 9/10$ in round t . Thus, in each round $t \in H'$, irrespective of prior rounds, we have $\Pr[X_t] \geq 81/100$. This implies: $E[X|\mathcal{E}] \geq (9^2/10^2)|H'| \geq \frac{9^2}{10^2} \cdot \frac{9^2}{10^2}|H|$ and

$$E[X] \geq (19/20) \cdot \frac{9^4}{10^4}|H| > 0.62|H|.$$

Just as in the proof of Lemma C.2, we can now apply Chernoff bound to obtain that $\Pr[X < 3|H|/5] < n^{-3}$ if γ is large enough. This completes the proof by recalling that all rounds of X are low contention and low interference rounds for node v . \square

F Proof of Thm. 5.1: Non-Spontaneous Broadcast

Theorem 5.1 *Assume the edge change rate τ is sufficiently small. There are constants $\beta, c > 0$, such that when running BCAST(β) in the synchronous mode, each node v receives the message in $O(D_{st}^c(s, v))$ rounds w.h.p.*

Proof. Let us fix constants $\rho, \hat{\eta}, \gamma$ so that Props. 3.2 and 3.3 hold with interference threshold $\hat{\mathcal{I}} = \min\{(1 - 1/\rho)^\zeta \mathcal{I}_c(\epsilon/2), \mathcal{I}_{ack}(\epsilon/2)\}/10$, where $\mathcal{I}_{ack}(\epsilon/2)$ and $\mathcal{I}_c(\epsilon/2)$ are the parameters corresponding to precision $\epsilon/2$. We set the passiveness parameter of TRY&ADJUST to $\beta = \gamma + 5$ and the stable distance parameter $c = 12\gamma$. Consider a node u that receives the message in round t . Let v be such that $v \in N(u, \epsilon)$ in the time interval $T = [t, t + c \log n]$ (and both stay alive during T). The theorem follows by an induction and union bound from the claim below. \square

Claim F.1. *Node v gets the message of u during T , with probability $1 - O(n^{-2})$.*

Proof. We prove the Claim by contradiction, and assume that v cannot get the message during T . Let us split T into phases (for u) and, similar to the proof of Thm. 4.1, classify the phases H into types: (type A) at least $1/10$ -th of rounds in H are high contention rounds for $u - P_t^\rho(u) \geq \eta$, and (type B) at least $9/10$ -th of rounds in H are low contention rounds for $u - P_t^\rho(u) < \eta$. We show that with probability $1 - O(n^{-3})$, all phases in T are of type B. Consider a type A phase H . We know from Prop. 3.2 that, with probability $1 - O(n^{-3})$, there is a set S of $\Omega(|H|)$ nodes in D_v^ρ that deliver their messages and restart TRY&ADJUST during H . We split S into three subsets, S_1 – those inserted by churn during H , S_2 – those inserted by edge change during H and S_3 – the rest.

Consider S_2 first. We know that in each local neighborhood, there are at most $\tau|H|$ nodes arriving due to edge changes. Note that D_u^ρ can be covered with $O(\rho^\lambda)$ local neighborhoods; hence, the total number of nodes arriving in D_u^ρ due to edge changes is $O(\rho^\lambda) \cdot \tau|H|$. Since $|S| = \Omega(|H|)$, setting $\tau = O(\rho^{-\lambda})$ a small enough constant gives $|S_1| \ll |S|$. Now consider S_3 . By the NTD primitive, for each node $u \in S_3$, all nodes in $D_u^{\epsilon/2}$ also restart TRY&ADJUST, setting their transmission probability to $n^{-\beta}$. By the setting of $\beta = \gamma + 5$, the probability that a node restarting TRY&ADJUST in phase H transmits again in H is at most n^{-5} . Thus, with probability at least $1 - O(n^{-3})$, the nodes in S_2 constitute an $\epsilon R/2$ -packing of D_v^ρ and are at most $|S_2| \leq O(\rho/\epsilon)^\lambda = O(1)$. As for S_1 , the probability that a newly arriving node transmits during phase H is at most n^{-5} , so $|S_1| = 0$ w.h.p. These bounds lead to a contradiction – $|S_1| + |S_2| + |S_3| < |S|$, if constant γ is large enough. Thus, each phase H of T is of type B with probability $1 - O(n^{-3})$. It remains to recall that T contains constant number of phases.

Now assume that H is of type B. In a similar way as in Theorem 4.1, we can apply Prop. 3.3 and show that if $p_t(u) = p_0$ at the beginning of H , then by the end of H , the effect on $p_t(v)$ is equivalent to applying at least $|H|/10$ of $p_{t+1}(v) = \min\{2p_t(v), 1/2\}$ operations on p_0 . Since $p_t(u)$ at the beginning of T is at least $n^{-\beta}$, it will take at most 11 phases to increase the probability to the value $1/2$. Then, a single phase will suffice for u to deliver its message (including to node v), with probability at least $1 - O(n^{-2})$. \square

G Spontaneous Broadcast

The basic observation is that if all the nodes start running BCAST* simultaneously, a *constant density dominating set* can be computed in $O(\log n)$ rounds (cf. [47]), where an r -dominating set of density κ is a set S of nodes such that for each node u , $1 \leq |\{v \in S : u \in N(v, r)\}| \leq \kappa$. The dominator algorithm is as follows: all nodes run BCAST* simultaneously, and 1. if a node v stops by SuccClear then it is a *dominator*, 2. if a node u stops by detecting NTD of node v then it is dominated by v . Having formed a constant density dominating set, it remains to disseminate the message using only dominators. In order for the set of dominators to derive the connectivity properties of the original graph, we need that (V, d) forms a $(\epsilon R/8, \lambda)$ -bounded independence *metric* space.

Theorem G.1. *In the static spontaneous setting, there is a uniform algorithm that performs broadcast in $O(D_G + \log n)$ rounds, w.h.p.*

The algorithm for spontaneous broadcast consists of two stages: 1. compute a constant density $\epsilon R/4$ -dominating set DS , 2. transmit the message using only the nodes in DS .

Dominating Set. The dominating set is constructed by running BCAST* in the spontaneous mode, i.e. all nodes start running the algorithm simultaneously. Let DS be the set of nodes that stop the algorithm by SuccClear. Since the dominated nodes use NTD($\epsilon/2$) as a stopping condition, DS is a $\epsilon R/4$ -dominating set. Moreover, DS is a $\epsilon R/8$ -packing, which implies that each in-ball D of radius $k \cdot R$ contains at most $O(k/\epsilon)^\lambda$ dominators. In particular each node u is dominated by at most constant number of dominators. The proof the algorithm terminates after at most $O(\log n)$ rounds is almost identical to the proof of Theorem 5.1.

Broadcast. Recall that we assume that the communication graph is connected and has diameter D_G . Also, since d is a metric, the graph is undirected.

The broadcast part is as follows: in the first round, the source node transmits the message to its neighbors; each dominator u , upon receiving the message, transmits it in each round with probability $p_t(u) = p_0$ until detecting $\text{ACK}(\epsilon/2)$, where p_0 is a small enough constant.

Note that as soon as all dominators successfully transmit at least once, all nodes will get the message. Moreover, if the constant p_0 is small enough, then the two algorithms can be run simultaneously: the key point is that the number of dominators in each cR -neighborhood is bounded by a constant (not depending on p_0). This ensures that the nodes *need not know* n , in order to coordinate the two algorithms. It remains to show that the message will get to all dominators in $O(D_G + \log n)$ rounds using the broadcast algorithm.

Consider a graph H defined over the dominating set DS , where for every pair of nodes $u, v \in DS$, u, v form an edge if $v \in N(u, \epsilon/2)$.

Claim G.2. *The diameter of H is at most D_G .*

Proof. Note that for every pair of nodes $u, v \in G$ with $u \in N(v, \epsilon)$, the corresponding dominators u' and v' are adjacent in H . Indeed, by the definition of the dominating set, we have

$$d(u', v') \leq d(u', u) + d(u, v) + d(v, v') \leq \epsilon R/4 + (1 - \epsilon)R + \epsilon R/4 = (1 - \epsilon/2)R.$$

Let u', v' be two arbitrary nodes in DS and let P be the path of length at most D_G in G connecting nodes u' and v' . By the observation above, if we replace each node $w \in P$ with its dominator w' , we obtain a path P' of length D in H , connecting u' and v' . This completes the proof of the claim. \square

Claim G.3. *For every transmitting node $w \in DS$, the probability that w delivers its message to all its neighbors in H is $\Omega(1)$ in each round t , if constant p_0 is sufficiently small.*

Proof. Let S denote the set of neighbors of node w in H , i.e. the nodes of DS that are at distance at most $(1 - \epsilon/2)R$ from w . It suffices to show that all nodes in S receive the message in a fixed round t where w transmits with $\Omega(1)$ probability. This event holds if the interference at w is no more than \mathcal{I}_c – event \mathcal{E}_1 , and no other node transmits in $D_w^{\rho_c}$ – event \mathcal{E}_2 , in round t .

By the properties of the dominating set, the contention in a ball of radius $\rho_c R$ is at most $O(p_0(\rho_c/\epsilon)^\lambda)$. Similarly to the proof of Prop. C.5, it can be shown that the expected interference by nodes in $\bar{D}_w^{\rho_c}$ is $O(p_0)$. Thus, by Markov inequality, $Pr(\mathcal{E}_1) = 1 - O(p_0) = \Omega(1)$ if p_0 is small enough. On the other hand, since the contention in $D_w^{\rho_c}$ is $O(p_0)$, $Pr(\mathcal{E}_2) \geq 4^{-O(p_0)}$ follows by Lemma B.2. Thus, $Pr(\mathcal{E}) \geq Pr(\mathcal{E}_1) \cdot Pr(\mathcal{E}_2) = \Omega(1)$. \square

Given that each informed node in DS has probability at least $\eta = \Omega(1)$ of successfully broadcasting the message, the rest of the proof essentially follows along the lines of the proof of [6, Lemma 6]. We present a sketch of the proof for completeness of the argument.

Let S_t be the set of nodes in DS that have been informed by round t , where R_0 contains only the source node. Let us fix a node $v \in DS$ and let d_t be the distance (in graph H) from v to the nearest node in S_t . Note that $d_0 \leq D_G$. The difference $\delta_t = d_{t-1} - d_t$ is the progress made in round t and is a Bernoulli random variable with $E[\delta_t] \geq \eta$ for all t when $d_t > 0$ and $\delta_t = 0$ otherwise. Let $\tilde{\delta}_t$ be a random variable that has the same distribution as δ_t when $d_t > 0$ and is an i.i.d. Bernoulli random variable with mean η otherwise. Let $\Delta_t = \sum_{i=1}^t \tilde{\delta}_i$. Note that node v has been informed by round t iff $\Delta_t \geq d_0$. Thus, we need to bound the probability $Pr(\Delta_t < d_0)$. Let $Z_t = \Delta_t - \eta t$. It is easy to show that the sequence Z_t is a submartingale. Moreover, for a round $t \geq \frac{c}{\eta}(D_G + \log n)$ and constant $c > 1$, $\Delta_t < d_0$ implies that $Z_t < -(c-1)(D_G + \log n)$. Since the sequence Z_t is a submartingale, we can apply Azuma-Hoeffding bound to show that $Pr(\Delta_t < d_0) \leq Pr(Z_t < -(c-1)(D_G + \log n)) < n^{-O(c)}$. Now the theorem follows by union bound over all nodes, by choosing the constant c large enough.

H Proof of Theorem 5.3: Necessity of NTD

Theorem 5.3 *For every (possibly randomized) broadcast algorithm \mathcal{A} that uses neither node coordinates nor NTD primitive, there is a $(\epsilon R/8, 1)$ -bounded-independence metric space where \mathcal{A} needs $\Omega(n)$ rounds to do broadcast in a $O(1)$ -broadcastable network, even if the nodes have CD and ACK primitives and operate spontaneously.*

Proof. Assume that $\zeta = 2$. We present the construction using a distance function d : it is then straightforward to construct the corresponding path-loss matrix. Denote $\delta = \epsilon/(8(1 - \epsilon))$. First, assume the nodes operate non-spontaneously, i.e. a non-source node may need to receive a message in order to start participating in a protocol. Recall that $R_B = (1 - \epsilon)R$. Consider n points p_1, p_2, \dots, p_n , such that for every $i, j \leq n - 2$, it holds that $d(p_i, p_j) = \delta R_B = \epsilon R/8$, $d(p_i, p_{n-1}) = \mu R_B < R$, $d(p_i, p_n) = (\mu + 1)R_B > R$ and $d(p_{n-1}, p_n) = R_B$, where $\mu = \epsilon(1 + \epsilon)/(1 - \epsilon) < 1$ (see the diagram in Fig. 1a). Clearly, this set of points forms a $(\epsilon R/8, 1)$ -bounded independence metric space. We place n wireless nodes at distinct points p_1, p_2, \dots, p_n uniformly at random. Let v_i be the node at p_i . We assume, further, that communication only happens according to SuccClear; in particular, if the interference at a node is more than \mathcal{I}_c then *none* of its neighbors receives its transmission.

Note that broadcast in this network can be completed in 2 steps, starting at any point. Note also that v_n cannot be directly reached from nodes v_i with $i \leq n - 2$. Moreover, if at least 3 nodes v_i with $i \leq n - 2$ transmit simultaneously, *no* node receives a message, including the potential communication between v_{n-1} and v_n . This follows by observing that $P/R_B^\zeta > \mathcal{I}_c$: the signal power at the neighboring nodes must be more than the interference threshold, otherwise a node could receive two signals of the same power.

Assume v_i with some $i \leq n - 2$ is the source. At the first round of the algorithm, nodes v_1, v_2, \dots, v_{n-1} receive the message. Let \mathcal{E}_t denote the event that v_{n-1} transmits and no more than $s = 3$ nodes v_i with $i \leq n - 2$ transmit in round t . By the observations above, v_n will not receive the message until \mathcal{E}_t happens. Moreover, during the subsequent steps of the algorithm, until \mathcal{E}_t happens for the first time, all nodes at v_1, v_2, \dots, v_{n-1} will have the same history and will be symmetric with respect to CD and ACK primitives.




Figure 1: Distance diagrams of the lower bound instances.

It remains to bound the expected time until \mathcal{E}_t happens for the first time. We can assume w.l.o.g. that in each round there are no more than s nodes transmitting. As discussed above, if v_n is not informed in such a round, then the nodes that transmitted in the given round can learn at best that none of them is at v_{n-1} and stop transmitting thereafter. Thus, we can assume that after each unsuccessful round there are at most $s + 1$ nodes that stop transmitting and the probability

\mathbb{P}_r of success at round $0 \leq r \leq n/(2s)$ (if no success occurred before round r) is at most

$$\mathbb{P}_r \leq \frac{\sum_{t=1}^{s-1} \binom{n-sr}{t}}{\sum_{t=1}^s \binom{n-sr}{t}} \leq \frac{(s-1) \cdot \binom{n-sr}{s-1}}{(s-1) \cdot \binom{n-sr}{s-1} + \binom{n-sr}{s}} \leq \frac{s^2}{n - (r+1)s + s^2} \leq 2s^2/n,$$

where the numerator counts the number subsets of an $n - sr$ -element set that contain a fixed element (the node v_{n-1}) and no more than s elements in total, and the denominator counts all subsets containing no more than s elements. A further straightforward calculation shows that the expected number of rounds until the first success is $\Omega(n)$ (see e.g. [12, Thm. 1]).

In the case when the node v_n might try to act without receiving a message, the instance above can be complemented as in Fig. 1b. A similar argument works in this case too. \square